WHITE PAPER

The optimization dilemma

Testing at the heart of your organization



dotmailer

"8% of email marketers spend no time on strategy and planning whilst 29% spend no time on optimization"

Source: Econsultancy's Email Marketing Census 2014



Contents

Level 1		Level 2			
Thought leadership		How to			
Why marketers don't test	4	Step 1: Making a case for test & learn	14-17	Reference	33
Why marketers should	5	Step 2: Setting goals, outcomes & learnings	18-21	About dotmailer	53-54
What should & can be tested	6-7	Step 3: Selecting the test elements	22-29		
How to create a test & learn culture	8-9	Step 4: Setting a structure	30-31		
Optimization & technology	10-11	Step 5: Measuring, interpreting & making	32		
Testing in the market	12-13				

Why marketers don't test

"The truth is, most marketers don't have enough time in the day to do all that is asked of them. All of us in the industry know it.

The most successful marketers find tools that enable them to quickly test and learn, improve and move forward continuously."



Skip Fidura.
Chairman of the DMA Email Marketing Council & Client Services Director at dotmailer.

This paper is concerned with testing impact rather than efficiency. The latter covers areas such as deliverability and rendering both of which are essential production processes however they are excluded from our focus.

One of the most important attributes of direct and digital marketing is that it is based on facts. It is provable and proven. Indeed this is one of the principles that all direct marketers learn.

All marketers aim for – and the best achieve – demonstrable results. They can show a start point and an improved end point; they can show the cost, the return and the reason. The discipline underlying this, and the mantra that many of us develop and repeat, is test and learn. It would be reasonable to assume that this process is the heart of programs and strategy. Reasonable yes – but wrong.

Econsultancy's 2014 email marketing census showed that 8% of email marketers spend no time on strategy and planning whilst 29% spend no time on optimization. Similarly, Marketing Sherpa's 2013 Email Benchmarking Survey revealed that 23% of CMO's considered A/B and multivariate testing to be somewhat or very unimportant in determining and communicating the value of the channel. As we all see on a daily basis, there is pressure to prove the value of the channel so it is, to say the least, odd that the very means of proving it is often not recognised.

The reasons why marketers don't test are pretty universal:

Lack of time – many email marketers are working to tight timescales that mean they are inevitably focused on production schedules. Setting up tests may involve new creative assets, new audience splits, additional sign-offs and certainly more planning; all of these take time which is already scarce.

Fear of distraction - often the schedule is driven by a business agenda such as sales targets for particular products and services. It's this that takes precedence and anything that may distract focus may simply be seen as a 'nice to have.'

Why marketers should

Cost control – it may seem perverse, but the relative inexpensiveness of email often encourages repetition and discourages adventure. It's simply easier to focus on other areas than assign thinking time to email.

Fear of uncertainty – combined with the cost issue above, where email can be relied upon to deliver results (it's a lever to be pulled at will if you like) there may be a temptation to avoid putting resources to fixing something that isn't broken.

Low expectations – combined with the reliability factor, often there is scepticism that testing can deliver any significant improvement that would justify the reduction in speed and increase in cost.

Fear of overheads – even if a business case could be made, there may be a perception that businesses lack the necessary focus and skills to create a measurable and robust test. Maybe this is based on the in-house resource, the agency, the ESP or any one of many. Maybe it's based on technology or people or both.

Concern about ability to change – Some companies may resist a test for fear that they won't be able to action what they learn. More often than not this is because of limitations of technology, process or culture, resulting in an inability to change. No one is suggesting that these reasons aren't valid. There are however some important and compelling arguments for making the leap. Businesses that have created a culture of ingrained testing have done so because:

Nobody's perfect – however effective your program looks, there is always room for improvement. An ongoing proactive testing process ensures that that this is realised. Industry and sector benchmarks will provide guidance on how marketers are doing against what others are achieving. The best marketers challenge themselves as much as they are challenged by others. They are champions and challengers at the same time. It's OK to fail sometimes along the way - not every test will win and companies that have created an effective testing culture and involve all of the right stakeholders, that mitigate the level of risk in all tests against any downside will succeed.

Recognition equals support – the board needs to recognize the value of a program and channel in order to invest in it. The best way to achieve this is to show a constant review and improvement in performance and value. A thorough and confident knowledge of the value of an email address and the ROI of a program will be gold dust in calls for justification and investment. Testing drives that knowledge.

It's about focus rather than resource – ESP software generally effectively supports testing. An increasing array of optimization technology is, similarly, generally well integrated with ESP's too. Effective testing therefore is unlikely to involve significant budget or disruption but rather it is about discipline and skills being at the heart of the process. This is largely about set up, objective setting, measurement and interpretation – they are core marketer's skills.

It's not just tactical; it's strategic – each test should act as steps towards achieving a business goal. Some very big advances have been made from what looked like very small insights. A lot of strategic thinking comes to reality when communications sit in the customer's inbox so that attention can be used and lessons can be learned.

Testing lives in the business not just in the marketing department – every business faces multiple challenges to which program tests can contribute answers. Marketers who sit side-by-side with colleagues in, for instance, merchandising, NPD or customer services can ensure that their email programs drive much bigger and more valuable insights.

It works and it's demonstrable – there will never be any doubt or argument about impact because, when properly structured, testing will deliver absolute answers and quickly.



What should & can be tested.

"Nothing is more effective at improving your email campaign metrics and ROI as split testing. And we're not just talking subject lines.

Done well, testing can double, or even treble or more, your open and click-through rates."



Michael Duxbury.
Head of Product Delivery, dotmailer.

With a commitment to testing, the humble email gives the marketer an awful lot of opportunity to do so. In this section we discuss some of the most popular example areas and what they can achieve (source - Marketing Sherpa's 2013 Email Benchmark Survey).

Subject lines - *tested by 86% of marketers* - it's the equivalent of a headline or the outer envelope in direct email terms; the attribute that has most impact on the recipient actually opening the email. A good subject line has been known to double open rates (source - email Mktg Reports).

Calls to action - *tested by 47%-62% of marketers* - the importance of size, number, positioning, wording and even colour can lead to uplifts of over 60% in clicks (source - email Mktg Reports).

Message/content/layout/images - *tested by 47%-58% of marketers* - the core message has a major impact on responsiveness once users have opened and viewed. Everything from the length of copy through to the size of images and how the message is expressed can matter. Uplifts of anywhere between 50% and 200% have been achieved (source - email Mktg Reports).

Time of day/day of the week sent - *tested by 46%-48% of marketers* - understanding when your recipients are willing to open and read your email is hugely important. Even relatively modest test have seen uplifts of both opens and clicks of 27% and 25% respectively. (source – Audience Point and Marketing Experiments Blog).

Landing pages - *tested by 44% of marketers* - if the email itself has effectively done its job then it has delivered people to a page at which the marketer aims to convert them so it's not surprising that working hard on that page pays dividends. Uplifts here can achieve 30%+ (source - our experience).

Target audience - *tested by 44% of marketers* - using segmentation (from general demographics to click behaviors) is often critical. Uplifts in clicks of 35% can be achieved; sometimes these can translate into even more dramatic uplifts of over 100% in resulting conversions (source – Jupiter research).

Personalization - *tested by 42% of marketers* - the ultimate aim of most marketers is matching content to user in order to create messages that are personal and not generic. At its most complex, this can mean creating segmentation systems and dynamic content strategies but there are simpler examples – from broad personas to simply matching product images to broad gender or age groups. The numbers explain why this is such an important area. Uplifts on open and click rates can reach 50% or even 215% respectively. This can follow right through to conversion uplifts of 175% (source - Socialware/Marketo).

From line/Sender - *tested by 32% of marketers* - Even an apparently simple area such as making the sender an individual rather than a brand can have a dramatic impact – uplifts in opens and clicks of 28% and 8% respectively are possible (source – Marketo).

Whilst these may be the most popular areas of testing, they are not the only ones. Each brand will potentially have a different priority and different areas of sensitivity in its messaging and communications therefore this creates a wide range of opportunities beyond these top 8 and including:

- Offer - At a generic level this may be part of the message/content or call to action categories; at a personal level it may be part of the personalization category above. It may however be a test in its own right given its importance as a key driver in response.
- Frequency – getting all of the other timing attributes (day of week and time of day) are important but if you don't know how often a customer wants to hear from you they may count for nothing. Fortunately this is something that can be tested.
- Format – we all assume that HTML (giving the opportunity for graphics beyond text) is most popular, but broken images will do you no favours here. Similarly, with circa 60% of emails now being opened on mobile devices it is essential that your emails are readable there – but how you best achieve this is another variable that can be tested. Video can have great impact but how it is delivered (e.g. by embedding or accessing via a click-through to a landing page) will need to be carefully considered and tested.
- Data collection – if you are trying to capture data from recipients in forms for instance, the length of the form and the types of questions that you are asking will be key to its success. Both of these elements can be tested; both will have a real bottom line impact on your campaign.

175%

Conversion uplifts possible, as a result of optimizing open and click rates.



How to create a test & learn culture

"We've learnt time after time that testing and learning must be ingrained and understood by the entire organization in order to succeed & continue to innovate."



Simon Bird
Chief Technical Officer, dotmailer.

The aforementioned examples illustrate that there's much to be tested and much to be gained from testing. As we have seen however, there are real challenges in doing it. It is important therefore to consider how those marketers who successfully adopted a 'Test and Learn' culture, made it happen. There are a number of fundamentals here that are essential to the case for commitment:

Making it recognised at the top table – in a world of 'Big Data' and narrowing gaps between MD and Marketing Director, visibility of the impact of digital marketing is being sought more widely. That's a good thing. Board investment decisions are always driven by hard numbers and that's what testing and learning delivers. A thorough and structured test plan should support the business and strategic marketing plan.

Making it part of the process – ensuring that test slots exist in every campaign will bring testing and learning to the heart of the process. This is nothing new. In old school offline DM it was often the case and yet, oddly, as the speed of campaign planning has increased, the discipline of using that to create learnings has simultaneously decreased. It ought to be the other way around. If the process is focused overly on production imperatives such as scheduling then the critical planning and evaluation disciplines of testing can be cast aside. They mustn't be. If testing is an essential element of campyaign set up then they won't be.

Making it visible and owned - enabling a wide range of stakeholders to contribute will ensure that testing is owned. That contribution may simply be a say in the areas to be tested or the distribution of results but making the learning an integral element recognised by a wide group will ensure that it is elevated beyond simply a tick box in the deployment process.

Making it scientific – in most cases, the outcome of a test will need to be planned and considered. It will rarely be taken from a single metric and will often need elements beyond the email itself to be evaluated (most commonly activity on landing pages and subsequent pages visited in click though journeys.) This means that a potentially wide group of stakeholders will need to be involved – from IT through Analysis/MI to Marketing and Planning – which is another reason for making it visible and owned.

Making it focused – testing for the sake of testing is a waste of time as is throwing variants up in the air when you aren't sure what they are intended to prove. It may seem obvious but frequently marketers are faced with so many variables (as can be seen from the testing opportunities above) that deciding what not to do may be a major challenge in itself. Successful marketers set up hypotheses to prove or disprove (e.g. that price reduction will drive click through's or that limited time offers will bring forward purchase intentions or that product images rather than packaging will generate better click rates etc) and focus on them. This will mean that all of the necessary work (creating assets, copy, audiences, tracking etc) will be focused on learning something specific and important. Nothing less makes sense.

Making it continuous – testing never ends. Today's champion is tomorrow's loser. Especially in markets where products or pricing change frequently, it is foolish to take a test as being any more than a measure of the time. Some principles will stand marketers in good stead over time but a proactive ongoing approach will be important.



"Testing never ends. Today's champion is tomorrow's loser."

Optimization & Technology

Huge potential. Simple execution.



"Recent advances in optimization technology are underlining the advantages of testing and creating exponentially larger ones.

We can deploy and measure tests quickly of course, but increasingly we can do much more within the tests themselves."



Chris Boddice,
Marketing Director, dotmailer.

The potential win for your business will be huge. Like many digital channels, email can be planned, tested and assessed very quickly. By combining the flexibility of testing with a reduction in the traditional test/evaluate/roll out cycle, marketers today have a compelling argument to 'Test and Learn' on three levels - test more, test quicker and benefit quicker.

Here are some examples that the smartest marketers are using today :

Automated A/B and multivariate testing – this enables marketers to set up tests, find a winner and then deploy to the remainder of the audience in an automated process. Some technology even enables this in real time so that, as soon as a winner is found, all subsequent deliveries can be switched to that content. This removes any delay between finding a winner and realising the benefit (source - 8 Seconds).

Time of day and day of Week – with increased intelligence now available, it's no longer necessary to rely on rule of thumb or generic answers to the perennial questions around the best time to send. With individual analysis of all customers and prospects, the optimum times can be found and tested via an automated process. Marketers see the benefits with little overhead (source - Audience Point).

Subject Line – having a clear idea or steer on the likely best option between several subject line variants removes the subjectivity from testing here. Using previous live experience, technology exists to give you that steer or maybe to remove the need to run the test in the first place. This is another great example of intelligence built up over time being used to help the marketer to improve their performance – today (source - Touchstone).

Real Time Content – whilst digital display advertising has embraced real time content, email is sometimes seen as a relatively static channel in a world where everything is dynamic. Not at all. Marketers can now present content appropriate to the recipient at the time that they open the email. This can be based on time or stock availability, location or even weather. However it is used, it provides the opportunity to transform time sensitive content and to take testing to a whole new level (source - Movable Ink).

Relevance and preferences – the marketer's goal is always relevance to the recipient. We all try to achieve this by looking at our customers' and prospects" interests, buying history, preferences & options and the way in which they engage with the emails that we send them. This becomes ever more complex. It also creates ever more variants. We could be getting close to the 'segment of one.' Marketers need to understand the impact and practicality of creating a vast number of variants which then are turned into vast ranges of content. How preferences are collected and how they are used are increasingly the subject of important tests.

Automation – timing, sequence and channel all have a key impact on email performance. We know this but we don't always know how. The increasing mix of channels in comms planning has placed greater focus on this. Marketers want to know how a solus email might compare with an email preceded by an SMS teaser or followed up by a phone call. They want to be able to test a sequence of emails over a period of time driven by an event (e.g. time since last bought.) They want to tie up apparently unrelated events (browsing on a site a week before sending an email.) The increase in automation that allows Multi Channel Management increasingly brings answers to these questions much closer to hand. These are essential areas because all marketers accept that the consumer (whether at home or in the office) now demands greater relevance and expects brands to be across the channels of their choice. Before any brand can race headlong into this fundamental change (with all of the business impact, resource and cost that it entails), they will need to understand how it works and where the ROI sits. It's a classic Test and Learn situation.

Optimization is important. As we said before, testing never stops and that's what optimization is all about. The discipline of setting up tests and then applying the learning will always create business challenges in time and resource so increasingly this needs to be as automated as possible – whilst anchored to solid planning focus and principles. Optimization technology can help in this very respect.

175%

Conversion uplifts possible, as a result of optimizing open and click rates.

Testing in the market

Today's trends.

Despite a range of obstacles facing marketers' adoption of 'Test & Learn', the smartest marketers are persisting and seeing results. We see great work from clients large and small. We see marketers challenging themselves and their business culture. We see testing as a discipline vying for ever more attention.

When we look at the wider picture we see why things are changing. The impact of testing is becoming increasingly recognised across the board. Traditionally the largest and the smallest companies have been the most loathe to embrace it and yet in 2014 they have both shown significant increases in uptake – 16% and 15% growth over 2013 respectively (source – Which Test Won State of Online Testing.) Whereas in the past the largest companies may have struggled for focus and the smallest for resource, both are recognizing what the mid-range companies have seen for some time – that properly managed Test & Learn creates bottom line results.

The scope of testing is also receiving more and appropriate focus. 80% of marketers track results beyond the initial click – this full funnel focus on the ultimate outcome delivers the full value but also requires the maximum commitment. The fact that this level is now precisely what is increasingly happening is a telling measure of progress.

This growing awareness is being used to address some of the obstacles faced by marketers in the past. Lack of resources (27%), lack of time (23%), lack of knowledge (11%) and lack of budget (9%) are all cited as reasons for not doing more or better testing (source – Which Test Won State of Online Testing) but most are becoming less compelling as advances in optimization, automation and planning techniques take center stage alongside results.



With the confluence of greater focus on bottom line results, growing interest in the value of data, new optimization technology and – most of all – real bottom line results being seen, now is the time to embrace Test & Learn. It's time to come out of the text books and into the center of digital marketing on a day to day basis.

How to

Implement Test & Learn.

This paper focuses on a practical ‘How to’ approach. It is intended to give brands a roadmap on this most important of areas and to enable marketers to both create and action Test & Learn initiatives.

Our focus is on five steps:

1. Making a case for Test & Learn.
2. Setting goals, outcomes & learnings.
3. Selecting the test elements.
4. Setting a structure.
5. Measuring, interpreting and making changes.

Each step is covered in stages that encompass both the principles and the advised practical actions. Test & Learn is a fundamental principle of Direct and Digital Marketing and yet it’s all too often overlooked in campaign planning and management.

Some brands fear that they lack the resources or skills to undertake a scientific program and to act on the learnings.













Step 1:
Making a case for Test & Learn

It will be essential to have buy-in to the time, resource and cost of testing. The ease of achieving this will largely be driven by the degree to which the value of the channel is recognised within the business.

We typically see two scenarios here:

The email program is not recognised for its contribution to the business - typically this is because it is either not of sufficiently high profile, its direct returns are relatively small by comparison to other channels or it has low priority in funding and resourcing. The latter is a classic chicken-and-egg situation which causes frustration for many marketers. The response to this situation has to be to:

- Highlight the returns of the existing program where they are not known. This should mean both directly attributable returns (in terms of conversions or traffic generation) but also an estimate of indirect returns (such as the 'halo' effect of sales to recipients of emails where the sales cannot be directly attributed to a click from the email.) Both of these areas require some analysis – and the latter some work with other areas of the business beyond digital marketing – but both will potentially reap valuable rewards.

- Highlight the value of each email address that you hold. This will be based on calculating the costs of collection versus the return in sales over time and it will begin to bridge the gap between a 'record' and a customer (because that's what an email address represents and we sometimes forget it.) There is a cost of collecting email addresses –whether from your own site inventory or from third parties – so it logically follows that you should be able to show the value that it generates. This can be measured both in aggregate – e.g. each email address held generates £X in net revenue or margin – and at individual level if you have the costs of recruitment and campaign versus the conversion revenue. With a solid idea of the cost and return of each email address and each program run to it, you will start to have a much clearer idea of where our improvement objectives should be set.

- Identify the return of comparable programs by relevant business sector or competitors. Benchmarks for email performance (if not for conversion) and for test uplifts exist in many areas (see Information Sources) and they are essential to be able to make a case where you think that you could do better. Nothing focuses the mind more than knowing that competitors are doing better than you or that there are opportunities for improvement that you are not taking. Make sure that you have signed up to your competitors' programs and are well aware of what they are doing and how they are doing it – this is essential context when benchmarking. Don't be limited only to your sector and your territory – there may be related organisations in other countries doing great work and, if you can show that this is having hard results, then you will be in a good position to make a case for trying new things.

- Identify the trends impacting your performance both inside and outside of the business. Internal issues such as site conversion performance or competitive pricing and service levels are clearly critical – they are unlikely to be restricted to an impact on the email channel but they are important when considering a case for testing because they will directly affect the likely results. External trends include general industry delivery, open and click rates, unsubscribe rates and levels of opt in for new subscribers. A good example is deliverability where your own performance versus benchmarked competitive sectors is important but so too is the overall trend for the future because you are planning ahead not just for today.

The email program is recognised for its contribution to the business. The focus here is much easier. No work is necessary to illustrate where you are starting from so, instead, the focus is on where you could go. The same work on benchmarking and trends applies as above but here the focus is much more on the potential uplifts and resourcing (see below).

With the above areas addressed, the core of the case to the business will be made like this:

- Starting point v ambition and rationale – the starting point is a statement of what your program looks like today versus comparable programs. The rationale for believing that your program can improve in line with – or beyond – the benchmarks that you have seen can be made here. This rationale is likely based on the comparable strength of your product/service, program design and offer etc. There is inevitably a degree of subjectivity here but it's the opportunity to focus on what you think can and will be achieved – to set your ambition. Be confident but be real; this process should be based on honesty and openness and if you can't make a convincing case to yourself then you won't make one to management.

- Hypotheses & Goals – The case needs to be made for where the tests will be undertaken, the learning created and the improvement achieved. This starts with hypotheses. Think about the areas of improvement and what you are trying to show. A more detailed form of this will be necessary in Step 2 but this is the place to show your high level intentions. It must give confidence to management that you have the focus to make it happen.

- Impact on operations & resource – It will be important to recognise the business impact in time and cost across areas such as:

  - People – it may well mean that the campaign team will have to spend more time on setting up test campaigns which will have an impact on the number that they are able to deploy. Similar people impacts will be potentially seen in 3 other areas – 1) the development of new creative assets, 2) the involvement of team members from outside the Marketing Dept (such as web analytics for instance) and c) your own time in creating the test plan as well as that of the planning team within your Agency or professional services steam within your ESP.

  - Technology costs – this may be in the form of additional functionality and modules within the ESP platform or maybe in test specific technologies. There may be an impact on CPM rates if some test cells are very small and are impacted by rate cards for lower volumes.

  - Business planning – if the hypotheses to be tested include areas of fundamental change to the business – e.g. the impact of free delivery offer or changes in service commitments – then it will be important to recognise the impact on the business in identifying, planning and measuring these. Creating a learning that cannot be applied is pointless so this needs to be built into the test plan right at the outset and it is unlikely that marketers can do this in isolation.

- Uplift – be realistic. Don't set targets that you don't believe are achievable; challenge yourself and ask yourself just how valuable will this be – is it all worth it? Think too about exactly how that uplift is going to be measured – opens, clicks, downloads, enquiries, purchases etc (again, more of this in Step 2.) In many ways this is the key point of the process so go for headline numbers that you believe in; balanced with the costs and impact above, they will illustrate to the business that this is a focused and realistic plan.

- Timescale - a realistic evaluation of when the tests can deliver will be important. This covers all of the impact areas above as well of course as an assessment of opportunity based on a calendar of forthcoming campaigns. It is not unusual to find a range of desired tests that would take months to achieve even if one went into every single campaign. This is where a realistic view of relative importance needs to be taken – can you prioritise those big win tests to go into the earliest opportunities whilst the others are phased over time? Remember too that prioritisation of what to test should be based on prioritisation of what learning could be actioned.

This approach need not be onerous and certainly need not create a significant process or paper trail. A concise approach to each of the points can deliver a tight argument with an absolute yes/no sign off decision. Greater details will come from step 2 and will feed into this. Remember that the more that you do these, the more they become second nature – not only to you in creating them but in management in approving them. It's about making Test & Learn central to marketing planning.

Step 2:
Setting goals, outcomes & learnings

These three elements are all part of the same ultimate objective but are not, as they may sometimes appear, the same thing. Instead:

- Goals - what you want to achieve (e.g. increase sales from visits)
- Outcomes - how you will measure (e.g. more conversions or increased order value from visits)
- Learnings - what you will take forward and use to effect change (e.g. we achieve the outcome by increasing the relevance of the landing page to the email CTA)

Sometimes, especially in straightforward tests, the three will be very simple to define and may merge together. However, especially for broader business impact areas, it can be very useful to think of them separately.

Importantly, this three level approach can be helpful in breaking down what you are trying to do. Many businesses have a range of performance targets in which the three levels are amalgamated so that focus is difficult, e.g.:

Sales targets – if everything is focused on a bottom line sale (or even margin of sale) then it can be difficult to break out the essential elements leading to it such as the performance of the conversion funnel and the enquiry rate in generating prospects to the top of that funnel. Both of these are critical to the overall Goal and Outcome.

Awareness levels – at the other extreme, a focus around something like awareness relies on a whole range of 'soft' measures from prompted awareness in research through to engagement rates (opens, clicks, read or forwards) none of which have a hard sales conversion. Context in how this is measured will be critical.

These two examples illustrate why a robust approach to setting targets is important. Sometimes it is very straightforward but other times it is essential and very worthwhile to put aside time to properly set these.

Beyond the obvious - improving open and click rates - think about goals, outcomes and learnings in the context of your existing or planned campaigns. There are several ways of looking at this that will help set practical targets:

Focus on what you know about your programs – most marketers have a strong feeling about the strengths and weaknesses of their programs. Maybe they know that their targeting is limited by a lack of data or that their content is too generic. Maybe they are concerned about the frequency of emails or the strength of offer. These are logical places to start. They use your gut feel as a marketer and that is very important but they also enable you to add some science.

Focus on practicality – in this regard a structure such as SMART can be useful. This is a framework for setting targets based on their being – Specific, Measurable, Achievable, Realistic and Time Based. This technique is not rigid and can be adapted to your business and your situation. That's one of its strengths. Think of the five areas as very straightforward disciplines. The process of applying these shouldn't be onerous – it can be easily and intuitively applied.

Focus on hypotheses – all tests should set out to prove or disprove something. If the goal is the big target then the hypothesis is your hunch (or expert opinion!) on how it will be achieved. Test hypotheses are very often the key area in which marketers apply their knowledge – it's where their brand, product, market and competitor intelligence comes together. Using a structure of 'by doing X we can achieve Y', we can apply the SMART principles from above as a very day-to-day discipline. Try running through some with the SMART process in mind. For example, if our goal was to generate greater traffic from offer emails, our outcome would be more visits to the offer pages or to the physical store so our learnings could be based on hypotheses such as:

- A more prominent CTA in our newsletter will increase clicks by 5-10%.
- Varying the offer content according to the engagement of the recipient can increase open rates by 5-15%.
- Always sending our weekend offer emails on a Friday afternoon isn't generating the maximum footfall possible.

These examples show a range – the first two are specific in quoting target uplifts whilst the third is more general because it is based more on a gut feel. All three though can be specifically measured and in a fixed time scale – the third one is the longest timescale (because it will need several weekends to see a picture) whereas the other two will see results within the standard click cycle of a email. We can set SMART objectives for each hypothesis; let's take the first one:

- Specific – it is about achieving an actual uplift in clicks (the most directly attributable outcome of an email)
- Measurable – with the relevant click tracking in place, the impact will be seen without ambiguity and likely with no additional effort. But don't take this for granted – many tests have come to grief when people assume that site visits and conversions can be tracked before finding that they can't – all you then know is that you have generated clicks but so what?
- Achievable – a 5-10% uplift should be possible based on previous experience or other evidence. You will know where you have created uplifts in the past – or where others have - to know if this is a realistic number.
- Realistic – increasing clicks is an appropriate target; it is not setting an expectation beyond that (for instance by assuming increases in conversion rates) since that is not in the remit of the email itself. The uplift can – and should - be ambitious but it must also be practical.
- Time related – you will know the response cycle of your campaigns (normally the majority of response is within 48 hours) and of your offers (stock clearance or weekend events will change this) so both can be applied to the uplift objective.

When considering the 'Achievable' part of hypotheses, you will likely have several starting points of your own. You may have seen one of your campaigns show a spike in performance that you think you can explain. You may have seen other departments or colleagues do things in



their emails that worked and you think can be applied. In both cases you will have numbers as your starting point – it won't be a case of 'this might work' but rather of 'this could achieve an X% improvement.'

This is important as we have seen in Step 1 because it drives the headline reasons for doing it and the impact against which costs are measured. If you don't have these numbers however, it can sometimes be a daunting thought to have to create them – what if the bar is set too high and you are potentially hanging yourself out to dry or what if they're too low and the business sees the proposal (or worse the results of it) as unambitious and pointless? Help is at hand. Your ESP provider can help with guidance and there are published sources as well as groups within social media in which email marketers share their experience and opinions. These can be invaluable. A number of examples are included in Section 4 to give you guidelines in the specific areas of testing but let's look at a couple here too:

An airline achieved a 97% uplift and 16% more site page views by using a 'Click Here for More Info' button versus a 'Find Out More' button - the former was of course larger (solely because of the word count) but in all other respects it was the same. Most impressively this Call To Action was right at the foot of the Newsletter and relatively unimpactful compared to the colour, images and copy of the body. Clearly, few brands would set a target of this sort of spectacular result but it should give you real confidence in what can be done and it should help you set your own level.

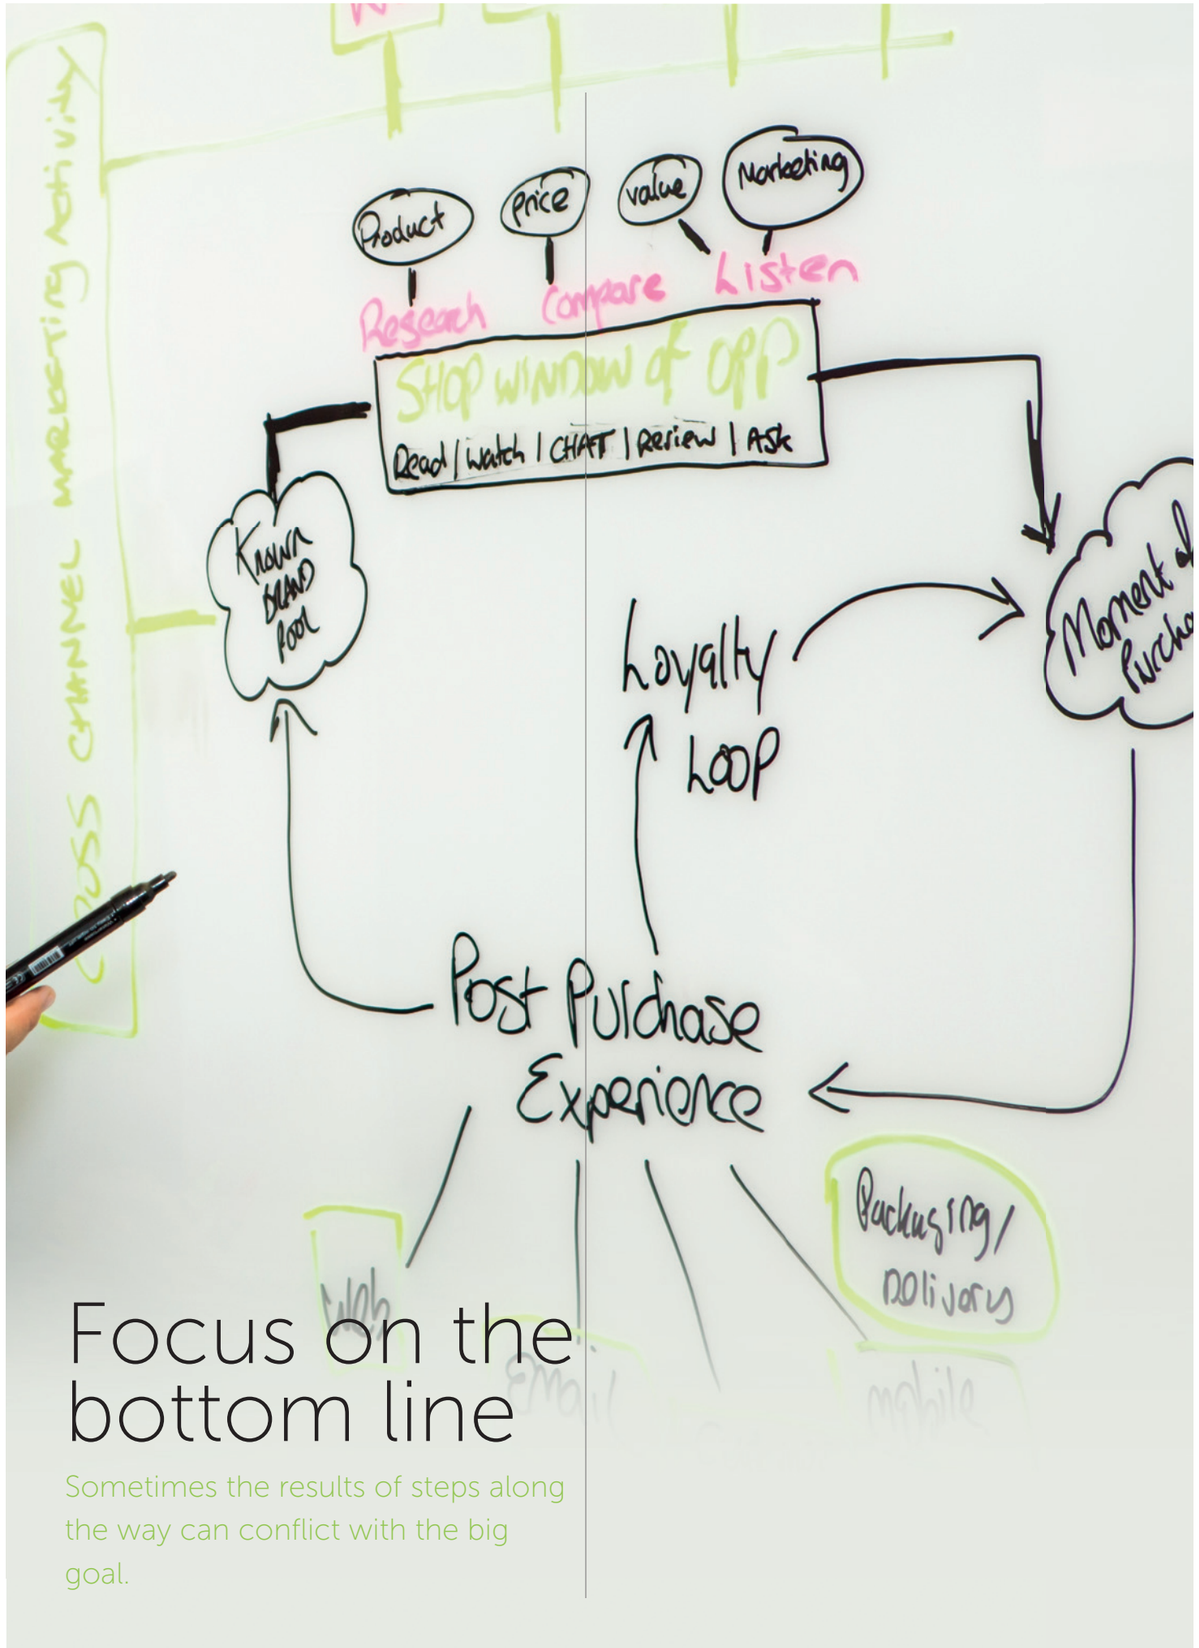
An insurance provider achieved uplifts on opens (5.8%), clicks (25.3%) and revenue (64.8%) by using a subject line focused on convenience rather than quality - all that simply form a subject line test! A few words focusing on one aspect of a product versus another. A simple hypothesis around which of the two angles would generate opens would have been a good starting point but this brand went much farther (and were able to because they measured each step of the journey.) It paid off spectacularly. Again, you may not choose to be so bullish on bottom line impact but this will give you context and confidence.

Of course you will need to apply your judgement and expertise in your own business and sector before applying other people's results but there are useful starting points out there for you.

Focus on the bottom line – sometimes the results of steps along the way can conflict with the big Goal. Focus on that Goal needs to remain at the heart. Take for example a test of frequency in which you are trying to measure whether a short copy weekly email might be better than a longer copy monthly one (this is a fairly typical issue.) Think about what 'better' means for your Goal. If your Goal is to generate sales then the opens and clicks of the email won't be the key metrics to measure. Indeed they may give you entirely the wrong result in that the monthly email may achieve a better engagement but the weekly one will generate more sales and maybe higher Average Order Values. This happens. Patting the team on the back for achieving increases in clicks won't make much sense when the overall sales achieved take a dive. This is why Goals are so important and why they must stay central.

Focus all of the key minds – Any actionable hypothesis has to be reached with all of the impacted stakeholders involved. In email terms that almost certainly means the site build and analytics team because they will be key to the delivery and measurement – e.g. it will be critical to have full conversion tracking from email to site if you are measuring beyond the click. Without these stakeholders involved during target setting you don't have a test.

So, setting Goals, Outcome and Learnings is a key discipline in successful testing. The good news is that you will have the majority of the information, instinct, intelligence and awareness at hand – inside your head! What is important to add to that is the focus and process to make it happen.



Focus on the bottom line

Sometimes the results of steps along the way can conflict with the big goal.

Step 3
Selecting the test elements

Marketers can feel that this step is akin to a kid in a sweet shop or, on the other hand, to someone drowning in a sea of endless possibilities. Both can be true. The beauty of email is that most things can be tested; that can be both exciting and daunting. The process of setting Goals, Objectives and Learnings as well as Hypotheses will fundamentally reduce the options but don't be surprised if you are still left with quite a range. That's a good thing.

The best test elements are really defined by what you are trying to achieve however it is worth casting the net wide because often the best outcomes can come from elements that:

- You didn't see as significant.
- You hadn't tested before.
- Aren't directly related to the goal and outcome on which you are focused.

The starting list of elements is very long but let's consider it and, in particular, the range of impacts that specific elements have been known to achieve.

Subject lines – as the first thing that a recipient sees, there's a strong case for making subject lines the starting point for testing. Combined with this fundamental response impact is the fact that they have a low resource and time impact too. It's important to remember that an awful lot can be tested – and achieved – from a small number of words used wisely. Good subject line tests can measure the impact of:

- Length. For example, a bold headline versus a more involved message
- Emphasis. Varied approaches such as value versus convenience position as in the insurance example in Step 2. Even within the line itself it is possible to measure the impact of one over the other. Similarly, think about an attention grabbing headline versus descriptive and informational headlines. The words themselves and their expression can be tested – for example, the use of capitals and key value words (conventional wisdom is that the word free in caps is an absolute non-starter because it will be blocked as spam but actually those days are gone. It is risky but not certain and it has been known to work.) Make sure that you use automatic spam scoring up front to help in this area
- Personalization. The inclusion of the recipient's name can often be good for a short term uplift but less consistent over time – the receipt can see them as gimmicky so beware.
- Symbols rather than words. As above, this may be good in the short term but not for a long term strategy.
- Brand over subject/offer. In combination with the sender field, your brand will likely be a major asset to you but sometimes it may detract from the offer that you are making.

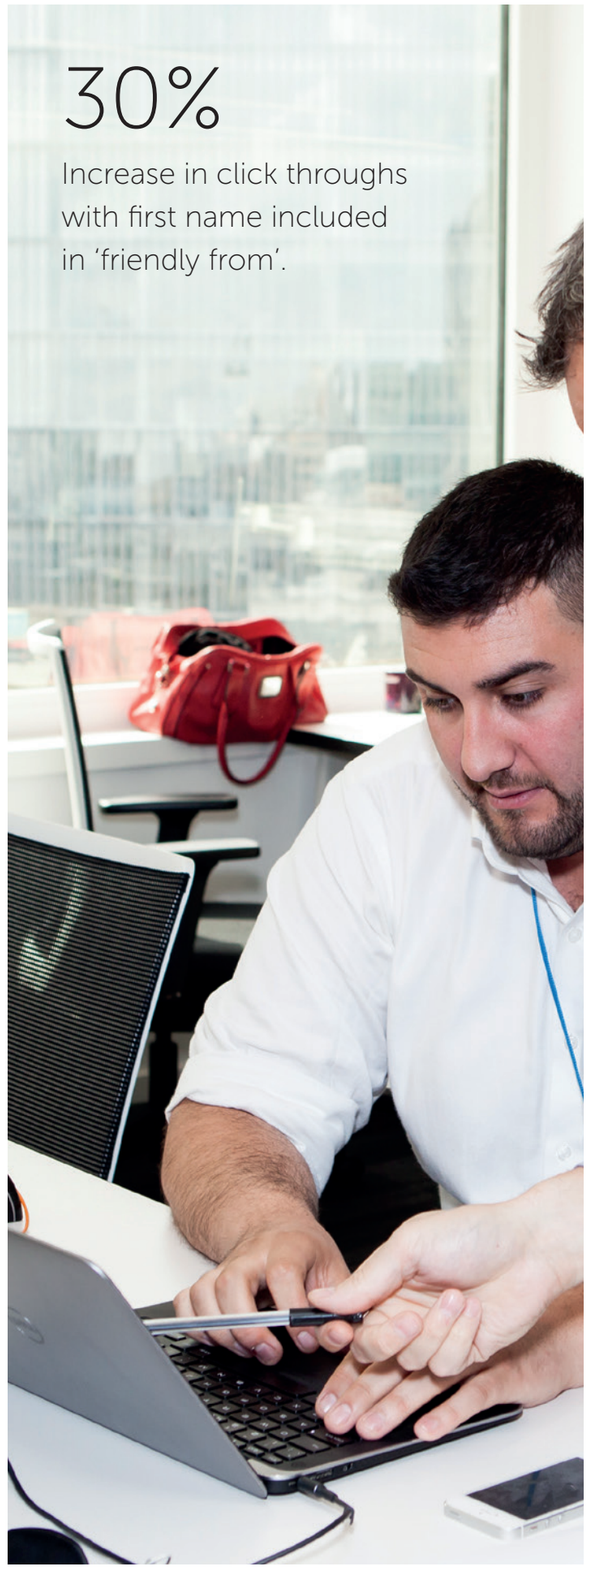
As we said in section two, it is essential to keep a close eye on what you are trying to achieve. It may be possible to create a high open rate by making headline grabbing promise in a subject line but that may just lead to disappointment when opened – the best tests maintain that uplift beyond the open.

From field / sender ID – whilst far less recognisably important than subject line, this is another area of major impact because it is one of the first things that the recipient will see. Your brand will of course be key in this but there are other aspects that can be tested most notably the use of a name (so rather than 'From Brand X', try 'From John at Brand X') which has been shown to produce notable uplifts in opens of 8% and click through's of over 30% *(Source Hubspot.)* A big return for a small change.

Personalization – personally addressing the email to the recipient sounds obvious but some brands are happy to treat an email almost like a door to door drop (a generic message without any acknowledgement that you know the recipient may well smack of little effort) or may lack confidence in the quality of their own data. The impact of several aspects can be tested:

- Salutation. Formal (using a title) versus informal (using a forename or even a nickname) versus maybe varying according to the communication itself.
- Copy or image. Personalization goes beyond the salutation and can be at the center of the email in reflecting the recipients' location or status of their most recent contact for instance. This will take more effort so you will want to prove its value.
- Offer. At the simplest level this can be positioning the offer as being relevant to that customer or group ('as one of our Top 10% of customers' etc) whereas at the most complex it may mean creating a different offer by recipient – see Offer and Dynamic Content below.

Audience/List Source – recognizing the audience, segment or source is clearly of great importance; in fact, it is often cited as the most important ROI tactic used by marketers. Just what this covers though can vary in sophistication. It starts at this basic level – long before producing dynamic content for each segment or individual – because there are many areas in which impact can be measured:



- Recognizing the audience's last interaction. For example, 'a lot has changed since you last visited us' or 'it's been a while since we saw you'. In advanced programs, brands will trigger emails based on the last content that the recipient clicked on (e.g. from their last email) or downloaded; they may also target by the type of content that the recipient has previously opened or read. This is great but much can be done even without such a level of sophistication.

- Triggering by an event. Rather than sending emails within a batch, a triggered approach can be a step on from the above. This way the content is varied by the audience and the audience is created by a Behavior. Brands that are able to instigate an email based on such Behavior (e.g. when a user has viewed site content, an email is sent within minutes reflecting that content) have seen real results such as uplifts on open rates of up to 65%, on click rates of over 45% and on conversion rates of 100-250% *(source – Jupiter Research.)*

- Recognizing the most immediate response event. Many brands send short term follow up emails to those who clicked (thereby expressing engagement) or those who didn't (to stimulate engagement.) This is relatively simple to do and will give you a good feel for whether or not it can work (see also Sequence and Frequency below)

- Matching products or services to audience interests. This can be at broad product group level - we don't need to have full segmentation and propensity models to know that certain product ranges will suit males rather than females or those living in cities rather than in rural areas. Even at this sort of level, it is possible to test the marginal impact of moving away from generic content for specific audiences by varying copy or images for instance.

- Recognizing personas. Broad audience 'personas' can be built on generalised attributes and limited to just a few in number; they can still deliver key learnings.

- Value of subscriber origin. You might except to see differences in quality according to permission sources recruited via your own site versus via third parties' or from a single opt-out to a double opt-in; this can be tested by measuring engagement right through to conversion so that you know unambiguously the actual value of each source for which you are paying.

- Recipients by their activity with you. Active customers are bound to view your emails differently from inactive enquirers so setting distinct targets and measuring accordingly makes good sense. Even with generic content you will see differences here and they may help make a case for a more custom approach

Dynamic Content – the next step beyond audience and list source testing, introduces far greater granularity. Here we can measure the impact of:

- Targeting an audience based on a segmentation system. This is likely created on the basis of multiple customer attributes (in effect, a step on from a broad persona or a general single variable such as time since last purchase; it makes use of combined known attributes such as customer source, buying activity, previous marketing responses etc.) The broad content will be varied by segment

- Creating individual content – the ultimate 'dynamic content' objective which can achieve a segment of one. With a flexible matrix approach we can set a range of different offers and messages to groups or individuals according to a range of attributes.

This approach makes the best use of email's flexibility and your own data. It can also deliver spectacularly – uplifts in opens of 50%, clicks of over 200% and conversion rate of 175% *(Source – Socialware ROI Study.)*

However, this takes time, commitment and budget; many brands are not yet there yet they can still achieve much without it via the Personalization and Audience List Source approaches above.

Send Time – one of the most frequently asked questions is 'when is the best time to send an email?' The answer will of course vary by recipient and maybe by type of email so the generalisations that are often quoted may be of little use. Testing – and intelligent analysis – can really help to measure the impact of:

- Varying by day (often incorporating weekday and weekend)

- Varying by time of day. Normally using broad bands for AM, PM, evening etc, this is relatively easy to do. Some marketers feel that sending very early means that the brand will grab the recipient's attention when they first open their inbox at the start of the day but beware those why might not take kindly to being alerted to an email when it arrives via their mobile in the middle of the night)

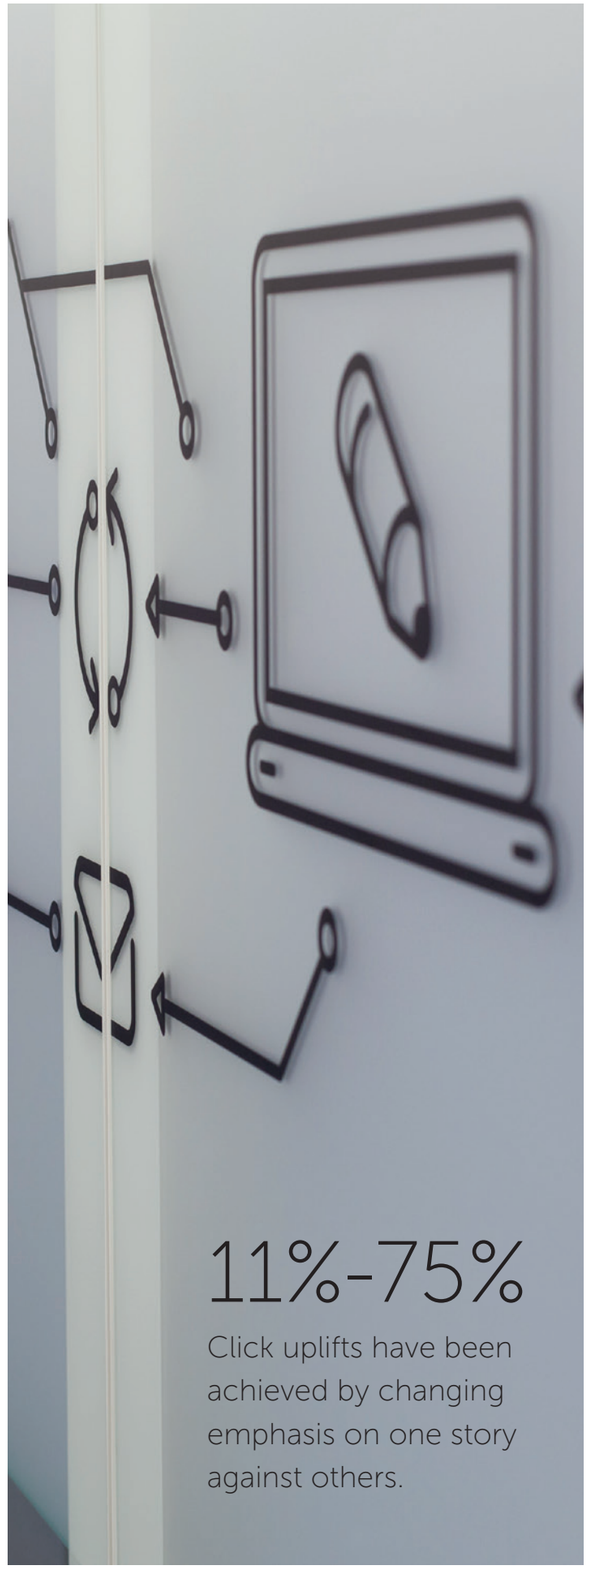
- Varying by proximity to an event. Typically retailers will choose a time to promote an in-store or online event according to when that event takes place. They know the notice that customers typically need or like.

- Varying by when the recipient is known to be online. Typically this might be based on when they originally opted in.

- Calculating the optimum time for each individual user. This is the ultimate in that, by analysing the history of a recipient's opens and clicks by time and within emailing type, it is possible to decide the optimum for that person. This requires a commitment to both a) the analysis and b) the 'throttling' of subsequent campaigns (so no more 'blasts' but more of a drip feed) but the results have been impressive with some brands achieving uplifts in opens of over 70%, click through's of over 150% and even in Average Order Value of 20% *(Source – Audience Point)*

Design/Layout – the appearance of a email is obviously essential in grabbing the recipient's attention. It needs to be right otherwise they may never get to your carefully crafted copy and beautifully optimised Call To Action. Like many attributes of a email, it is often surprising what big differences can be made by apparently small tweaks. Think about the key areas of impact that can be measured:

- Emphasis. Most emails have more than one thing to say especially in Newsletters. Defining your 'lead'



message or the overall structure will be hugely important in grabbing attention – and the right attention. The use of 'hero' images and stories is well proven but it may not work for everyone. Testing is important here because wins can be big. Click uplifts of 11%-75% have been known by changing emphasis on one story against the others; even 5% uplift in converted sales have been achieved by this. *(Source – Which Test Won cell H6 and H19)*

- Format. A message can be expressed as block copy, bullet points or video. There are advocates of all (including many who believe that the advent of video means the death of words) but the actual answer for you and your recipients is yours and theirs alone. So test. Results have shown uplifts of as much as 30% here *(Source – Which Test Won cell H10)*
- Balance. Any email is a mix of copy and images but the range of combinations can have a major bearing on how the whole is seen. It is very likely that any marketer may want to show their product but do they show it alone, in the context of their customers' use or in its packaging etc? Alongside the image how do they describe it in order to best underline (rather than detract from) the image? This is the area of balance that has generated some marketers click uplifts of more than 75% so it is well worth the effort. *(Source – Which Test Won cell H21)*

Copy – this is obviously a critical area because it's where you need to best get your point across. Every word and comma can count. Aside from the essential message itself (which you will be clear on) think in terms of the three key areas of impact to measure:

- Length. Long v short is one of the oldest debates. It's also one of the areas with the most established – and frequently broken – rules.
- Tone of Voice. You have a brand voice and this is the place to use it but you will also have different tones that you may want to use according to the message and the emailing type (for instance a service message versus a sales message.) Slight variances in tone can make a difference
- Emphasis. There are of course many ways to say the same thing and significant tests of 'focused sell' versus informational news are fairly typical test areas which have been known to spring some significant surprises.

Offer – whilst many email campaigns may start with a generic offer there are ways in which this can be flexed. There are good reasons to try because clearly the offer is at the very heart of the proposition so it could well have a major impact. It is possible to measure impact in:

- Varying the expression of an offer. Typically, retailers might offer free Postage and Packing on items of a certain value which, in effect, is the same as offering a £ or percentage discount on the item itself. Where these expressions have been tested against each other they have been seen to make a real difference – conversion rates increasing by over 60% *(Source – 8 Seconds)*
- Varying the offer itself. Maybe some recipients are more worthy of a better or one more appropriate offer to their specific history. This is where offer and Personalization overlap. If a case can be made for varying the offer by recipient just remember the practicalities which will be important for measurement. If the most loyal customers are to receive a better offer than the newer and less active customers you will want to track the revenue that it generates and the likelihood of it being incremental versus the fear of your giving away margin to people who would have bought anyway; you'll also need to ensure that the offer cannot be redeemed by anyone outside of your favoured group. As a consequence of these points, this is a more involved area of testing.

Call to action – the impact of a CTA can be quite dramatic. That's not really surprising because, if we have done our job so far, the recipient should be sufficiently motivated to progress and the CTA is how we get them to make that next step. Testing can measure the impact of:

- Location. Making the ask must be at the right place. Not too hidden so that the recipient is frustrated but not too prominent that they haven't yet read the proposition and therefore don't have a reason to respond.
- Wording. Sometimes a straight 'Buy Now' is all that's necessary but often it's not as straightforward as that so 'Click Here For More Information' or 'Find Out More' are amongst the many variants here. Setting the recipient's expectation is really important. It matters too – click through uplifts of 50% are not uncommon when you get it right. *(Source – Mark Brownlow, e-mail Marketing Reports.)*
- Colour and design. The CTA has to jump out at the recipient so its design will play a big part. Images and words both matter here so consider testing both – separately and together.
- Channel. If you have a CTA for a specific channel then consider using smart technology that delivers a response to the easiest place to convert. For example if you are promoting a mobile app then giving an iPhone user a CTA that goes directly to the app in the Apple App Store will make life much easier for the recipient than simply giving them CTA's for Apple, Android and Windows and leaving them to make the decision. This very example has produced uplifts in clicks of over 230% *(source – Movable Ink.)*

Often the impact of each of these individual points can be significant but in combination they can be transformational. Uplifts have been seen of 115% to over 300%. *(Source – Which Test Won cell H14 and H23.)*

Frequency – 'how often should we send a email?' is a perennial question asked by marketers keen not to tire their audience but also not to miss opportunities to sell. We all know just how clogged up our inboxes become (the stats tell us even if we can't see it with our own eyes) so it's a legitimate question. It's also a legitimate area for testing because that is likely the only way that you are going to arrive at an answer. It's not really a single question either:

- Depending on the type of email - (say a Newsletter versus a product update) the answer will vary and sometimes it is also reliant on a combination of factors. Brands who have tested short copy on



Frequency

'How often should we send a email?' is a perennial question.

frequent emails versus long copy on infrequent emails have found that the answers are not one dimensional. Some emailings are more susceptible to timing than others

- Take for instance abandoned carts - where the conventional wisdom of sending a single email as soon after the event as possible is regularly challenged by brands sending hours or even days later and over several steps rather one – sometimes with significantly better results.

Sequence – many brands now employ multi touch programs especially where their Call to Action is not a straightforward sale. The 'drip feed' effect can be greater than a single message. Understanding how these perform however is a classic case for testing. Measure the impact of:

- Time. The period between each touch often matters in keeping people engaged but avoiding making them tune out
- Roles. Think about which message is the teaser, the main communication and the follow up.
- Channels. The sequence above does not need to be solely via email and, in fact, it rarely is. Think about email and SMS for instance where the roles of teaser and follow up can be tested and even interchanged.

Landing page navigation / user experience – all of the hard work that has gone into creating an impactful email can be wasted if the post click experience falls apart. Sadly, often it does. Maybe this critical gap is missed because the responsibilities lie with different teams - that's one of the reasons that we need to 'focus all the key minds' in step 2. With the email design and desired outcome in mind however, there are many important areas of impact to measure:

- The principles. Copy, layout, CTA as above – they are as important at the destination as they are in the departure point (which is what your email is.)
- Focus. The landing page needs to deliver on the message that you have expressed in the email. If it doesn't then you risk 'dumping' the recipient in a place where they can't easily do the very thing that you have motivated them to do. This happens.
- Urgency. This should be the last step to the recipient converting (be that buying, downloading, registering etc) so waste no opportunity to stress the NOW. In combination with the above Focus, this cannot be overlooked.

User control – much email marketing is 'push' based. We marketers decide what to say to people and when to say it. We have their permission to communicate but beyond that the decisions are ours. Switching that to a situation in which the recipient makes decisions - i.e. where they 'pull' - can be scary but it can also pay dividends. There are good reasons for this. Many brands pay too little attention to engagement rates and too much to aggregate campaign results. They know that the campaigns are generating acceptable bottom line numbers and that responses are on schedule but they don't know that vast percentages of their audience are rarely, if ever, even opening their emails let alone responding. Many people opt out of their attention long before they opt out of their permission. Giving a recipient the opportunity to make more choices as an alternative to unsubscribing – which bands are increasingly doing – seems a bit like locking the stable door after the horse has bolted; why not test the impact of that choice much sooner? There may be a fear of giving choices lest people decide to receive less but an unengaged audience is not giving you attention and without attention your marketing is going nowhere.

There are two clear areas of impact that can be measured:

- Preferences. Measuring how many people take up the opportunity will tell you the appetite
- Outcome. Measuring the response and conversion from emails sent according to the recipients' preferences rather than the marketers schedule will be the ultimate call.

So we have outlined 14 broad areas and a number of flavours within each. That's a lot of testing but not all will be relevant to you. It's worth the time to think about areas that could have the biggest impact and to start there. Strike a balance too between:

- Big Contrasts such as two complete design variants. Sometimes these headlines can change everything but they're risky and demand resources and time.
- Small Changes such as CTA positions. Sometimes these can have a disproportionate affect but other times they can give you minor results that make little real difference to what you will do going forward.

It's hard to get the balance right when there are so many choices. There will be successes and failures but that's the inevitability of testing. Most important, a good focused plan will deliver more winners than losers.



Copy, layout, CTA

They are as important at the destination as they are in the departure point.

Step 4
Setting a structure

With your test elements chosen and your hypotheses in place, it will be important to set up a structure that will deliver a result. This is sometimes thought of as the 'stato step' and it's true that there are statistical rules that underpin it but these can be applied relatively simply. What is more important is that the foundation is in place to ensure that you have an actionable and compelling result. That takes marketing nous as well as numbers.

There are five key points here:

1. Decide on how you will test.

In setting test elements within your hypotheses, you must ensure that when you see a difference you know that it came from a specific element and that such a difference matters. That's where structure comes in. Generally speaking this means one of two routes:

- A/B Tests – these literally split an audience in half and show one variant to one half and another to the other. It's the purest and simplest method. The key discipline of course is ensuring that everything except the chosen variant that you are measuring is the same – from the audience through timing and content to post click experience.
- Multivariate (or MVT) Tests – these measure multiple variants and their combined impact. In effect, MVT aggregates the combination of elements into a result. Needless to say, this is far less a manual process than A/B and requires software to do the 'heavy lifting.' Theoretically there is little limit to the number of elements and combinations but in reality (as with all testing) you need to keep focused on what it is that you are aiming to prove and the size of audience necessary to do it (see below.) As with A/B too, it is essential that all other elements are equal.

There are pros and cons for both. A/B tests are generally quicker and easier to setup and they're more intuitive but they are of course limiting in scope. MVT removes that limitation but takes longer to plan, requires many more variable assets from copy to images etc and requires a much larger audience. For these reasons, many marketers start with A/B and migrate to MVT when the number of tests and element variables warrant it.

2. Decide on what you are measuring the test against.

It is hugely important that your test elements are compared against a properly 'ring fenced' comparison. In general there are two methods generally used for this:

- 'Control' which can be either a) the known status quo/current best practice version against which new variant elements are to be measured, or b) if there is no current best then a representative sample of the audience to be held back from the test and served with something alternative for comparison. Without having this control group in place, it is impossible to measure a comparable uplift and any results will lack context.
- 'Fallow' which represents a group of customers who receive no marketing comms at all. The buying or engagement Behavior of this group is the purest lightning rod for the overall impact of marketing activity rather than a measure of specific campaigns. Care needs to be taken here though in selecting, refreshing and matching the expectation of the users. In the latter case for instance, if people have signed up to an email program based on explicit promises of content they cannot be kept in a fallow group otherwise you are not meeting your commitments to them. It sounds silly but you might be surprised.

3. Decide on the significance level that you need.

To be sure that the results of your tests are reliable and can be used to make business decisions, you need confidence. In statistical terms there are levels of confidence that can be applied when setting up a test. Most tests are set at 95%+ which in effect means that there is only a 1 in 20 chance that the results are not representative of the overall population – i.e. you can be 95% confident that, had you run the same campaign against your total audience rather than the test audience, you would have seen the same result. This level of confidence is high and there are arguments that marketers can work at far lower levels; this is important because the higher the level of confidence the greater the audience size so it may well have an impact if your audience is limited. There is no right or wrong answer to where you set the confidence level – it's a business decision. There are easy ways of setting the level (see Information Sources) which simply require you to specify the current rates and the expected/ desired uplift (both of which should be part of setting the hypothesis covered in step 2); you don't have to be a statistician.

4. Decide on the size of audience that you need.

This is important because you will want to avoid exposing too many of your subscribers to a test that may not work (and yourself to the risk of poor performance) but likewise you will want to ensure that the audience size is large enough to be significant for the elements that you want to measure. The significance issues above come into play here for the latter point. Tools to calculate the minimum sample size are available (see Information sources) based on level of confidence that you choose. This will give you a guide but not necessarily the full answer – if the minimum recommended test cell size is much smaller than your available audience then you may decide to go further depending on the level of importance that you place on the test and your own gut feel as to its importance and likelihood of success. Again, the statistical tools are there to support you but they're just a foundation for your own marketing expertise and ambition.

5. Decide on the purity of test.

This may seem odd; surely, we would want absolute purity for the most accepted result wouldn't we? Real life isn't quite so simple of course. Some things create 'noise' that can't be avoided. Here are four examples:

- Seasonality - Whilst no marketer would run a test in December and consider it representative of the year overall, there are many other less dramatic examples of where time of year will impact on results.
- Time - Obviously the ideal is to run all versions of a test against each other at exactly the same time but sometimes this isn't possible (e.g. a welcome email might be sent 24 hours after a registration but that means that it emails on different days according to when each registration happens.)
- Site Changes - You would never run email tests that deliver traffic to a site that itself is undergoing changes. Sometimes however scheduled changes are unavoidable and, even if they're small, they will reduce the purity of the tests.
- Duration - Can you tell what you need from a single emailing (especially where the test outcome is reliant on delivering site traffic) or might you need to have several days' worth of site experience before making such a call? Tools are available to estimate the duration of a test (see Information Sources) based on site visits and likely conversion and these can be invaluable when setting expectations.

Each of these are examples of context – you will need to consider them and to make decisions based on the likely impact that you think they will have on results. Like all elements of testing, it will be about a good balance of disciplined planning and pragmatic acceptance.

The 'implementation step' is really important but, in many ways, it is the most automated and process driven. It makes best sense as a part (albeit an integral one) of the overall plan. The plan matters.

Step 5
Measuring, interpreting & making changes

If you have set your goals and structure then this final step should be the least involved.

You will be confident that:

- You are testing the right things and the right number of them
- They can be properly measured through the funnel
- You have built in allowances for volumes & issues such as seasonality
- You know when to declare a winner (for instance it may not be the response after 48 hours but the conversions after a week)
- Everyone has bought into the exercise.

Everyone will therefore be waiting to see what has been learned.

The key points then are:

- Share the results. If you have bought together all of the stakeholders in planning then it's essential that they are bought together again in reviewing
- Keep it simple. Present the result (either in person or in written form) as very clear Q&A. Recount the hypothesis and then outline how the results have proved or disproved it. You don't need to go into the confidence level or statistical backgrounds as to how the test samples and control were created (that's just detail); it's the headlines that matter. Try to take ownership of this yourself to stress that you started the process and you have been there throughout. If you have technical and statistical support at hand that's great but focus the discussion around the marketing imperatives.
- Focus on 'so what.' The big issue is what to do next especially where the results signal a significant change going forward.
- Agree the next steps. Presenting the results isn't the end. The 'So What' should signal it as the beginning. Many effective test plans fall away at this most critical of stages because there is no clear path to next actions. It would be naïve to assume that every stakeholder will 100% buy into the results – there will be debate and there may be dissent – but there should be enough consensus to create a plan for what will be done to action the learnings.

Remember that a test result isn't forever and isn't necessarily for all! Winners in one program may not be in another and, in any event, they don't stay that way over time. Keep this as the heart of your thinking because testing – and more importantly learning – never stops.

Information sources

The following is a range of sources and resources that we think you will find useful in devising and running your Test & Learn programs.

Making a case for Test & Learn

- UK email Benchmarking – DMA or dM equivalent
- UK email Census – DMA and eConsultancy

Setting goals – outcomes & learnings

- SMART principles - http://www.projectsmart.co.uk/10-steps-to-setting-smart-objectives.php

Setting a structure

- Minimum sample size calculator (inc confidence level options) - http://www.rrdonnelley.com/RMS/ServicesResources/MinSample.asp
- Duration test calculator - https://vwo.com/ab-split-test-duration/

About dotmailer

Our mission is to make it as easy as possible for marketers to get results that make a dramatic improvement to their business.

Email marketing. Evolved

dotmailer is the UK's number 1, email marketing automation provider.

In 2009, we developed a successful, innovative ecommerce platform, used by flourishing retail brands like Rodial Rush Hair, Nicol Farhi and Myla.

Today, all our ecommerce innovation, experience and expertise goes into enabling maarketers and online retailers to do the kind of advanced email marketing they only imagined. And they're seeing someremarkable results.

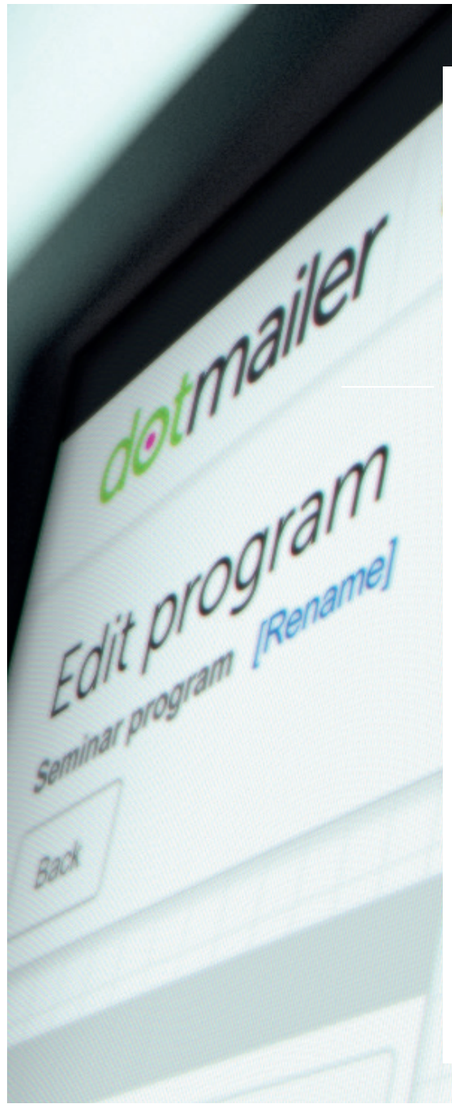
For more information and a free 30-day trial visit dotmailer.com, or call 0845 337 9170 to speak to an advisor.

Integrate with dotmailer
Just some of the platforms we have pre-built integrations for. And we're adding new ones all the time.

"dotmailer allows us to do things we have only dreamed of before."

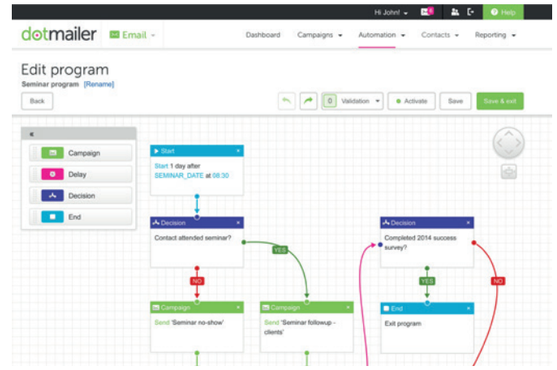
Amberley Sherman.
Marketing Manager at Cazenove + Lloyd





dotmailer's marketing automation. The automatic choice.

dotmailer lets you set up data-driven conditions that optimize, test and adapt content and creative based on campaign results, so you can send the right email, to the right person, at the right time.



Program builder at a glance

Starts - Establish the criteria for a contact to join your program with the start node.

Campaigns - In just a few clicks, you can select and send a campaign from the list.

Delays - Select the amount of days you'd like set for any follow-up campaign you'd like to send.

Decisions - Within the decision controls there can be a rule, or set of rules. It's entirely up to you.

Discover how dotmailer can help transform your business:

www.dotmailer.com

+44 845 337 9170

HEAD OFFICE
No. 1 London Bridge
London
SE1 9BG
+44 (0)845 337 9170

EDINBURGH
9-10 St Andrews Square
Edinburgh
EH2 2AF
+44 (0)131 718 6037

USA
350 7th Ave, Suite 307
New York, NY 10001
1-212-971-9408

CROYDON
No. 1 Croydon
Addiscombe Road
Croydon
CR0 0XT
+44 (0)845 337 9170

MANCHESTER
Pall Mall Court
61-67 King Street
Manchester
M2 4PD
+44 (0)161 618 1070

dotmailer

IMAGINATION...DELIVERED

Written by: Nick Fuller (Company Name as well?)