

Designing for Conversion

How to effectively optimise your online conversion rates.

by Richard Muscat

(Updated for 2019)

Contents

[A Tailor On A Hill](#)

[What is 'conversion' and how do I design for it?](#)

[Step 1: Give the most important answer](#)

[Step 2: Retain their curiosity](#)

[Step 3: Remove usability obstacles](#)

[Step 4: Understand the consequences of being wrong](#)

[Step 5: Limit choice](#)

[Why People Fail at Online CRO](#)

[Fail 1: Using metrics as a shield](#)

[Fail 2: Having the wrong motivations](#)

[Fail 3: Incorrect or split focus](#)

[Fail 4: Starting at the bottom](#)

[Five Rules to Get Things Right](#)

[Rule 1: There can be only one](#)

[Rule 2: Have data you can trust](#)

[Rule 3: Play to your lead time](#)

[Rule 4: Copy is key. Copy is key. Copy is key.](#)

[Rule 5: Aim for substance over significance](#)

[Calls to Action](#)

A Tailor On A Hill

This guide is about the craft of designing online sales – what many people have come to refer to as Conversion Rate Optimisation (CRO for short). The term encapsulates a range of work done by various kinds of web professionals aimed at increasing (optimising) the percentage (rate) at which visitors to a web site or landing page make a purchase (conversion) before leaving. This is a line of work you should be interested in if you run high-traffic ecommerce sites, depend on the web for lead generation or sell most kinds of software online. If this is you it's likely you could be doing more to optimise your sales, revenue, and customer lifetime and that's where CRO tends to fit in.

CRO is not a new topic and I regularly meet marketing people, graphic designers, UX specialists, consultants and start-ups who think they should be doing this work. The surprising thing for me – and main motivation for writing this – is how many seem to misunderstand the basic concepts or worry that it's an unethical line of work. So whether you're wondering if CRO will be valuable to your business, you've found yourself losing heart in it, or are weighing up whether doing it means selling your designer's soul to the devil, then hopefully this guide will shed some light on what you could do differently.

In practice the work revolves around marketers, developers, and designers testing out different ways of selling stuff online for the purposes of improving, or optimising, sales and revenue. If you're doing this work already, or are considering it, you've probably come across terms like iterative testing, split testing, multivariate testing, cohort testing, statistical significance, funnel analysis, landing page design, call to action optimisation and many others.

This book is not about those things.

I'm not writing about *The Secrets of High Converting Landing Pages* or giving you the *Top Ten Designs that Double Your Conversion Rate* or describing the *Seven Characteristics of Highly Successful Split Tests*. Partly because there's a lot of books and blogs already dedicated to these topics. But mostly because, in my experience, such "best practices" are often red herrings, particularly to newcomers.

The reason I call them red herrings is not because such best practice recommendations are incorrect. It's because by their nature best practices are generalisations; averages if you will. Besides the fact that they may not apply to your specific circumstances, they will, at best, make you average. And I'm assuming that as a designer your aim is probably to be at least a little bit above average.

So what I aim to give you instead is what you need to get right before you start testing anything at all. My objective is to help you work towards creating the right mindset for yourself, understand the underlying concepts and show you how to create your own set of best practices.

In reality, "conversion rate optimisation" is nothing but a euphemism for "sales". Many designers however seem to feel an innate sort of distrust of anything having to do with the S-word. I've heard many reasons for this which I broadly summarise into these four groups:

- *Authentically great design should need no sales! – If I do a great job people will tell their colleagues and friends and sales will happen organically.*
- *Talking to sales people makes me feel icky; they're always so pushy and coercive – I don't want to do that to my customers!*
- *I'm a craftsman and my job is to make the world a better place by producing high quality work – people will recognise and appreciate that and get more pleasure out of my work if I put 100% of my effort into it.*

- *Online sales is just so full of dark design patterns – I’m scared I’ll get sucked down a path of bait-and-switch techniques, dodgy affiliate marketing and pop-up ads!*

These are all valid concerns; ones I’ve struggled with myself many times. The story that I always fall back on to work through these conflicts is one I remember from my childhood days:

I grew up in a small village on top of a hill on a small island in the Mediterranean. My grandfather, *nannu* in our language, was the village tailor and for most of his life earned his keep from custom tailoring for the people of his village and beyond. He took enormous pride in his craft, making sure everybody knew he did all the work from scratch: designing, measuring, cutting, tailoring, fittings and alterations.

What he took most pride in however, the thing that would first trip out of his mouth when he arrived home, was making a sale. He was so proud of this precisely because he succeeded at doing something that he considered not to be his core skill. And he sold lots.

Spending time in his shop as a little boy one summer I marvelled how day after day, a never ending line of people would walk in, have quick chat, and almost invariably leave having paid for something, placed an order or made an eager promise to come back. At the end of one particular day, when literally every walk-in bought something or other, I had to satisfy my curiosity: “But *nannu*, how do you do it? How do you get all these people to buy?”

I refuted his initial claims that it was done by magic. Eventually he explained:

“These people” he said, “well, I’ve been working on persuading them to buy for years. Nobody walks in ready to pay – even if they came with the best word of mouth recommendation. At the end of the day nobody really wants to give away their money if they can help it – suits are

expensive and they don't have much money after all. The best I can hope for is that when they walk in they are curious. At that point I try to answer their questions but also keep them curious rather than try to get their money. Even if at that point I'm working on a very delicate job I stop and talk to them. But the most important thing of all is to have patience. You can't let them feel you're rushing them."

At this point you're probably thinking that this is starting to sound too much like an over-romanticised description of a wise old craftsman on a hill. And you're right up to a point.

Yet there's a core of truth in what the old man is saying. As designers – or tailors – we love our job because we believe that by virtue of our work people will have better experiences: be it a more intuitive bit of software or a more comfortable pair of trousers. If we truly believe in the quality and value of our work then dedicating time and effort towards selling that value – converting people – is something we should not only embrace, but take pride in. So always *design to sell*.

And the key to doing it without selling your soul is to have patience. Understand that converting customers, keeping them curious all the while answering their questions, takes time. Dark design patterns (using design to trick people rather than convince them) take root only when we become impatient and try to make a quick buck without aiming for true conversion.

What is ‘conversion’ and how do I design for it?

So what is conversion? Like many terms that transition to common usage their true meanings and origins can become diluted or forgotten. ‘Conversion’ as applied to business is not a new term. Sales and advertising people have been using it for decades.

In practical terms it means generating a sale: turning a prospect into a customer via the exchange of cash in return for a product or service.

In aspirational terms it should mark the point at which a convinced decision is made: an interested party is converted into a loyal and long-term customer. Think of it like making a meringue: once you’ve converted egg whites and sugar into a delicious confection you can never reverse the process. Similarly, converting prospects into customers should mean that they’re so bought into what they’re paying for they will never turn back (e.g. by asking for refunds, not purchasing again, complaining about you on Twitter, etc.).

Thinking about CRO in this way makes a big difference to how you practice it. First of all you begin to understand that conversion is a long-term process. It’s long-term along two dimensions: first it takes time for you to understand your customers well; and second it may take time for your customer to convert (exactly how long depends on what you’re selling). The crux is convincing and compelling persuasion not trickery and fast-talking.

The reason I said earlier that best practices can be red herrings is that a lot of them don’t speak to why a person converts. It is absolutely true that a well-placed call to action (e.g. a “Buy now” button) can make a big difference to your conversion rate.

As does its colour and size. Something that stands out is more likely to get noticed than an obscure text link.

However that's not *why* people convert.

Very few people will decide to pay you \$49.99 because you have a stylish and obvious purchase button. *They pay you that \$49.99 because you've persuaded them to click that button.* Conversion in fact happens before your visitor clicks that button. As a designer setting out to achieve conversions therefore there are five steps you should take:

Step 1: Give the most important answer

Understand the most pressing question on your visitor's mind and answer it immediately and without delay. This builds their trust in you and makes them feel their search is bearing fruit.

Visitors don't convert into customers *because* you are offering a free trial of your software. Visitors convert because you persuade them that your product is going to make their lives better. You need to take the time to find out what "better" means for the majority of your audience in a tangible way.

If you're selling products or services to other businesses, they're likely to be looking for ways of using your offering to increase their margins, acquire more customers, reduce their costs, delight their own customers, operate more efficiently, etc.

If you're selling directly to consumers they're probably looking for different things such as having better personal relationships, more earning potential, saving time, security for their kids, comfort, entertainment, etc. On the other hand they might be looking to run away from or mitigate against a perceived threat such as bankruptcy, guilt, missing out on opportunities or loneliness.

There's no formula for knowing which of these things your customers are looking for but there are ways of finding out. The best approach is to talk directly to customers on the phone or face-to-face and apply the "5 Whys Technique" (refer to the last section for this) to uncover their motivations for visiting your website. As you build up an 'inventory' of reasons for why they're visiting your shop, try to keep a tally of which of these various driving forces are most *frequently and intensely experienced* by your customers. The ones at the top of the list are the ones you should address first in your design.

If you can't (or don't want to) talk to customers directly, in-page surveys (such as those provided by Qualaroo or PollDaddy) are also effective. Ask your visitors very straightforward questions that give you good qualitative data to work with such as "If this page convinced you to try our product can you tell us why?" Asking vague yes/no questions like "Are you satisfied with this page?" is not going to yield rich answers that you can work with.

Step 2: Retain their curiosity

Even as you answer their most pressing question, plant a new question in your visitors' minds. This keeps them interested in what you have to say.

Your answer to your visitor's most pressing question should naturally and automatically give rise to a further question.

Let's say your product promises to save your customer some money by compressing their data backups more effectively thus reducing spend on expensive storage space. You've answered the most pressing question which was to explain the product benefit and how it works in terms the visitor can understand. It might be natural for your visitor to wonder what this means in practice to their business. You can lead them to the next step by asking something like "How many terabytes of data do you

currently manage?” The intended effect being to start them thinking down the path of how much money they could be saving in their own specific circumstances. You might answer this question by showing benchmark figures or by getting them to fill in their system’s details in a ‘calculator’ and providing a realistic number together with pricing.

Their next curiosity might be to find out how to sell this to their bosses if they’re not the decision maker. A “Find out how our customers convince their bosses” checklist might do the trick...

What I’m describing here are examples based on some projects I’ve worked on but they shouldn’t be what you go out and implement. In each case you need to find these things out for yourself from your own specific set of customers or prospects.

The takeaway point is that with every question you answer you should be ready to provide the next bit of information with the objective being that at some point in the journey the next most pressing question becomes “How do I pay for this?”

One way of doing this is to use the benefit/advantage/tactic approach. This works something like this:

- For most things you’re selling, you should be able to articulate one or more **benefits** that your customer will experience from purchasing your product. These tend to be overarching goals. For a fitness app, say, an overarching benefit might be a healthier lifestyle, more self-confidence or increased strength and stamina.
- Each benefit in turn gives the purchaser some **advantages** either over competitors or simply over their previous way of doing things. In our fitness app example, the corresponding advantages might be less anxiety, losing weight or being more likely to win a race.

- Finally **tactics** are ways of delivering those benefits and advantages. Typically these translate into product features. For the benefit of gaining self-confidence by losing weight you might offer tactics such as weight-loss focused training plans, alerts, reminders and peer support.

Your message and design will be more compelling if you manage to weave a story that includes these three aspects of your product. Where you start from however will probably be dependent on what your prospective customers are searching for.

Step 3: Remove usability obstacles

Ensure that the mechanisms for satisfying their curiosities are as usable and as obvious as possible. This makes sure their experience is not a frustrating one.

At each step of the way you are asking your visitor to take an action. This might be using the ‘calculator’ I described above, filling in a form, checking out or simply taking in a page of text and clicking “Next”. It is vital that at each stage in the customer’s journey you make the mechanism for proceeding to the next step and answering the next question as easy and as obvious as possible.

Your job here is to remove as many barriers as possible that may get in the way of your visitor proceeding to the next step and the techniques for doing this are derived from good old-fashioned usability engineering practices. Make sure each page has one clear and well-defined call to action. Make sure it’s easily discoverable and the text describes the action well. Make help text available where applicable and make sure forms follow standard design patterns.

Remember that ‘usability’ and ‘removing barriers’ do not necessarily equate to ‘less steps’. It is incredibly easy to assume that the less steps you have the better your conversion rate will be but that is not necessarily the case. Designing the length and steps of your customer journey is highly dependent on your product and audience.

Step 4: Understand the consequences of being wrong

Understand deeply what the consequences of a wrong purchasing decision are to your prospective customer. Don't rush to a sale without addressing them.

The length you design your customer's journey to be depends highly on the perceived consequences in your customer's mind if they were to make the wrong decision. The perceived consequences of being wrong have a big effect on how quickly you can convert prospects into customers.

Certain products and services can be purchased on so-called 'impulse' because the consequence of making the wrong purchase is low. The consequences of renting the wrong movie from iTunes for example are that you're out of pocket for a couple of dollars and your evening's hopes for cinematic entertainment go unsatisfied. At worst you and your spouse or your friends will have a laugh over what a horrible movie choice you've made and get a different one. So as a designer selling low-value products that are purchased on impulse you can probably benefit from short customer journeys with as little friction and as few steps as possible. Very often these are products for which both the problem to be solved ("I'm bored right now") and the ideal solution ("I want to watch a movie" or "I want to play a game") are very clearly understood by the customer. On-demand movie rental services, book stores and app stores can take this approach and minimise the journey as much as possible very often aiming for one-click purchases and rentals.

Other products may require longer customer journeys if the perceived consequences of being wrong are much more significant. Higher-value products or products solving complex problems can have bigger negative consequences in the event of a wrong choice. The most obvious negative consequence is related to price. A business or individual with a spending power of, say, \$100 dollars, will think longer and harder

about purchasing an item costing \$100 dollars than he or she will about purchasing an item costing \$1. Often more importantly however are the perceived consequences of being wrong about the benefits the purchase will bring and the opportunity cost of a wrong decision there. A business buying an expensive backup solution will be worried not only about the cash outlay but about what happens if the service doesn't work. The consequences to the business of choosing the wrong backup software are loss of intellectual property, competitive damage and, in the case of the individual making the purchase, possible unemployment. So the evaluation and purchase process here is unlikely to ever be a 'one-click' purchase. Trying to do that will make the prospect feel they're being rushed and are likely to lose trust in you.

In similar fashion to keeping a tally of motivations driving your customers, try to keep a tally of what your customers perceive to be the consequences of a wrong purchase. If you're familiar with project management techniques this is not dissimilar to keeping a 'risk log' with two dimensions represented: likelihood and impact. Use this tally to bubble to the top the items that have *both a high likelihood of occurring and the highest impact* in your customer's mind. You can then design your customer journey, layouts and messages to address these concerns in a way that's step with your potential customer.

Step 5: Limit choice

Too much choice can be debilitating. Give your prospective customers just enough choices so they feel in control, but not so many they feel confused.

In 2000 Sheena Ingear of Columbia University and Mark Lepper of Stanford University published a now famous paper entitled: *When Choice is Demotivating: Can One Desire Too Much of a Good Thing?*. In it they describe the outcomes of three experiments they ran to determine the effect of choice on the likelihood of buying and subsequent satisfaction with the purchase. The paper is worth reading for

yourself but in summary they showed that people are significantly more likely to purchase when presented with a limited set of choices compared to a larger set.

In one of their studies they offered a group of people a choice of 6 jams and offered another group of similar size a choice of 24 jams. While more people showed interest in the latter selection (60% compared to 40% for the smaller selection) the subsequent purchasing behaviour is astounding. Of those encountering the limited choice selection, 30% chose to make a purchase. Of those encountering the extended selection, just 3% chose to purchase. Their other two studies, involving gourmet chocolates and elective class essay assignments showed very similar results.

The lesson for designers and CRO practitioners here is to not automatically fall into the trap of assuming that more choice will lead to more conversions. Sometimes too much choice can lead to indecision or, worse, a decision to purchase that only results in a persistent nagging doubt that one of the other options could have been a better choice.

Why People Fail at Online CRO

At this point it may be clear that conversion optimisation work is not something you do on one Friday afternoon. It is continuous and iterative work that should be ongoing. Not appreciating this can mean that first-timers go into the field with unreasonable expectations which are unlikely to be met.

In both my in-house and client work I've noticed practitioners and organisations failing at successfully optimising their conversion rates because they're making one or more of the following four mistakes.

Fail 1: Using metrics as a shield

Using numbers meaninglessly or as a means of distancing customers rather than gaining intimacy. Numbers without conversation are like pasta without a sauce.

One of the wonderful things about the web, which has probably contributed to the rise of CRO as a web discipline more than anything else, is the availability of real-time and detailed metrics of website visitor activity.

If you have the right tools installed and hooked up to your website or web app you can literally watch visitors navigate through your website in real-time. You can watch traffic flow through your website funnel and change tack as visitors respond to your changing content and designs.

The danger of having such detailed information is over-reliance. One symptom of over-reliance is 'forgetting' to talk users and customers in person. It is significantly easier to open a web browser and check Google Analytics than it is to organise phone calls with customers or read through pages of qualitative written feedback.

The thing to keep at the forefront of your mind is that metrics are there to help you validate hypotheses you formulate in the 'real world' and not the other way around.

For example, working on Red Gate's shopping cart last summer, the metrics showed us that visitors were dropping out significantly at a particular point in the process. We put in a lot of work to simplify the number of steps on the assumption that the process was too complicated. The net effect was: no difference. Talking to customers directly we then found out that our visitors were not "dropping out". They were simply going through the process up to the point at which they got a price for their chosen products and add-ons and then going away to clear it with their purchasing managers. This often took a few weeks. No amount of simplifying was ever going to make them purchase there and then!

Another symptom of metrics over-reliance is focusing on the numbers when the numbers are meaningless because they are too small. In one particular project recently we had a development team basing product decisions on split test results where there were, at a given point in time, literally five customers using this brand new product. Four users preferring a particular feature over one is not significant in real terms even if it is statistically so. The right approach in this case should have been to have in-depth conversations with all of those five individuals.

Fail 2: Having the wrong motivations

Not understanding why you're doing it, including not having the right expectations for what sort of results you should be expecting.

One of my mentors, a telecoms consultant, once told me the secret to his success is constantly asking his clients an often irritating question: *"Remind me why we're we doing this again?"*

As a well-funded business, bootstrapped startup, designer, or consultant you have many options open to you when it comes to peddling your (or your clients') wares and increasing online sales. A full focus on web-based conversion rate optimisation is but one of them – and it's not always the right one. If you're not sure why you're engaging in CRO work it is very likely you'll be disappointed with the results.

The following scenarios are some examples of when CRO work is likely destined to fail or not produce a worthwhile return on investment:

- Doing it because it seems that “everyone's doing it” and unless you do it too then your web strategy *must* be broken.
- Attempting to make up for loss of market share and/or customer uptake due to your product not delivering user value.
- Engaging in CRO work when you have very low volumes of traffic and market reach with the expectation that CRO work itself will increase your brand's awareness.

The motivation for engaging in CRO work should be born out of a clear business need or goal that builds on existing product/market fit. The following are some examples of business goals that align well with CRO work:

- Fine-tuning your marketing collateral and landing pages so that you can roll them out across various web properties or make them available to affiliate networks.
- Getting the best return on investment possible on marketing campaigns.
- Significantly increasing your company's revenue by selling higher value products (upsell or cross-sell) or add-on services to a captive audience.

- Addressing a big and undesirable change in user behaviour at some point in the customer journey such as a drop out in a checkout or evaluation process.

These examples are by no means exhaustive. The takeaway is that whatever you're doing should be progressing business goals and have a specific target.

Fail 3: Incorrect or split focus

Focusing on irrelevant metrics, having a severely split focus or not even knowing what the conversion rate is.

It may seem like a no-brainer in hindsight yet one of the things most newcomers to the field get wrong is focusing on the wrong conversion rate. Or worse, something that is not a conversion rate at all.

First off, as a rule of thumb your business or project should only have one metric that is considered to be the true conversion rate. This may be different at different points in time but it should always be the one that has the biggest effect on the long-term success of your business. For example, if your start-up is in pre-launch phase you should focus on people signing up for notification when the product is released. If you're selling downloadable software your conversion rate shouldn't be sign-ups or downloads, it should probably be unit sales less refunds. If you're trying to validate whether a new product is valuable to customers you should probably be focusing on usage.

Examples of common mistakes are things like:

- Focusing on download and install rates when the focus should be sales.
- Focusing on advert click rates when the focus should be sign-ups
- Focusing on sign-ups when the focus should be paid subscriptions.

- Focusing on sales when the focus should be renewals.
- Focusing on sales without taking into account refunds.

The often hidden problem with these sorts of mistakes is that they lead to undesirable behaviour in yourself, many times without even noticing. For example, it is extremely easy to boost download numbers of a product by using 'bait and switch' techniques, focusing heavily on the fact that something is "FREE!!!" or implementing auto-download pages. You will boost your download numbers for a certainty but when you then look at your sales numbers you will be disappointed.

Fail 4: Starting at the bottom

Running optimisation tests too early, applying the wrong design patterns or not having the time (or people) to do it well.

The fourth most common mistake I notice is people starting CRO work by diving right in and setting up some complex split, cohort or multivariate test on a landing page or shopping cart. This is a bit like landing in Heathrow and trying to sell combine harvesters by engaging in door-to-door sales.

As an example, in 2012 I set up a small team that was tasked with optimising various parts of a corporate website. To show some immediate 'progress' we jumped straight into setting up a few split tests on the home page and a couple of other high-traffic pages. The tests were exceptionally disappointing because although we got significant results, we then didn't know what to follow up with. Although we proved that certain button placements resulted in more clicks we didn't have a solid plan for what to do next with that behaviour. Conversely, we later focused on a specific part of the checkout process that was upselling a support package. We did this after a few weeks of research by talking to customers and analysing historical data. By comparing the opt-in rate for the upsell in telephone sales to online sales we came to

the conclusion that there was substantial room for improvement and so we prepared a two-month test plan. After two iterations we achieved a 30% increase in opt-ins by changing the message from one focused on financial discounts to one focused around the benefits and advantages derived from premium support.

The actual testing stage of new designs, copy and workflows should happen at the point where you have one or more pretty strong hypotheses for your experiments based on your research. You've either identified a particular spot in your workflow or customer journey that is severely underperforming or else you've decided to target a particular area for improvement.

In either case you should have a pretty good idea for why your customers or users are behaving in a particular way and a reasonably high degree of confidence that you can change that behaviour by your proposed experiments.

Five Rules to Get Things Right

I'm proposing here five rules of thumb for successful conversion rate optimisation on the web. You don't need to follow all of these to get your optimisation work right and they're by no means exhaustive. But if you're following neither, you may want to re-evaluate your strategy.

Rule 1: There can be only one

As a business you should have only one metric that at a given point in time you consider to be representing your conversion rate. It is vital to have only one ultimate goal that as a business unit you're aiming for as this makes it easy to get alignment across the different business functions or between you and your client. In most cases this should be unit sales or paid subscriptions.

These are my rules of thumb for choosing a good conversion rate metric:

- It should represent the thing that you most care about as a business.
- It should always be a verb (purchasing, subscribing, renewing, using, etc.)
- It should be measurable with a good degree of accuracy

Finally, bear in mind that as a company or product changes over time what is most important to the business may also change. So it is a good idea to review your conversion rate metric periodically.

Rule 2: Have data you can trust

I would estimate that in my client and full-time work, half the time I spent on CRO involved debating the accuracy, believability and validity of the data we were looking at. *“What’s the margin of error on these metrics?” “What if this was just a seasonal effect?” “Do we know how many visitors have javascript turned off?” “Are you sure that’s the correct error margin?” “Did we have some downtime there?” “Did you hear that this industry leading analyst said that Google Analytics uses sub-optimal sampling algorithms?” “I’m not sure I believe that error margin is correct!”*

Having data you can believe in is of paramount importance. Aside from saving you from endless soul-destroying conversations, optimisation is all about improvement and to be able to demonstrate that improvement that you need believable data so that you can compare the before and to the after. If I or my client don’t believe the validity of the data we’re looking at it becomes very hard to know for certain what effect my CRO work is having.

These are my rules of thumb for having data you can trust:

- **Start collecting data from day zero**

The further back your history stretches the easier it will be to detect regular patterns, seasonal dips, errors and anomalies.

- **Outsource your data collection**

Unless you have a highly competent in-house business intelligence team use a third-party tool for data collection such as Google Analytics, HubSpot, KISS Metrics, MixPanel or whatever is most suited/affordable to you. While it’s tempting – especially for engineers – to want to roll your own, the operational cost and time needed for maintaining home-grown systems quickly balloons; and if you can’t keep up with it the data’s accuracy, and consequently its believability, diminishes.

- **Aim for consistency not perfection**

Perfect data doesn't exist despite data scientists' claims. There's always going to be a margin of error or a better sampling algorithm you could be using. But if you have the same error margin over the course of years then you don't have to worry about it. By being consistent with your choice of data provider (e.g. always use Google Analytics [or insert your favourite provider here]) then you can assume your error margin is going to be consistent and thus ignore it.

- **Keep an event log**

Stability doesn't really exist: there will always be things like system down time, traffic spikes, big marketing campaigns, unexpected press coverage or what have you that can skew your pretty and steady baseline numbers. The best solution here is to keep track of these events on a timeline so that, when you're looking at your graphs, you can see whether anything unusual happened during that period. Google Analytics for example provides a simple way of annotating your data timeline with notes that serves this purpose really well.

Rule 3: Play to your lead time

"Lead time" refers to how long it takes for a prospect to become a customer from first contact to the point at which they convert. Depending on the product or service you are offering this may vary from a few months to a few seconds.

Broadly and very generally speaking people go through three stages before they become your customers: 1) awareness of the problem to be solved and of your existence, 2) evaluation of your product, 3) decision to purchase (or not).

Understanding how long it takes for your prospects to go through these stages affects

both how you approach your designs and how long you need to wait before you can observe results.

In the iTunes movie-rental example I mentioned earlier, a customer might go through these stages in a matter of seconds. They're already aware of the problem and the product and they're impatient to solve it. Their time window for evaluating a movie choice will be in the range of seconds to a few minutes, perhaps by looking at star ratings, reviews and the movie trailer. If you make changes to way you present movies to customers you could expect to see the effect of those changes within a few days if not hours.

On the other hand, one of the products we sell at Red Gate helps software developers identify and fix problems with their applications' memory management. Over the years we've consistently found that a large number of developers are simply not aware that this is a problem at all, let alone one worth solving. So a lot of the effort needs to be put into upfront education that may take weeks or months to take root and the effect of changes to the customer journey may only be observable months down the line.

In fact this is one of the places where those best practices I mentioned in the introduction can lead you down the wrong path.

Rule 4: Copy is key. Copy is key. Copy is key.

'Copy' is text. Technically it refers to the words that appear in an advert or on a web page. I like to use it more broadly as a way of referring to writing down in actual words what the message of a design – be it an advert, landing page, home page or shopping cart – should be.

In my early days of working as a graphic designer I hated words. They got in the way of my beautifully crafted illustrations. They never fit the grid and would justify in the most horrible of ways. Later, as I turned my attention to UI and web design, I formed the opinion that users would only read and need text if the UI wasn't easy enough to use. And when I first engaged in CRO work at Uniblue in 2006 I remember running several split tests that seemed to indicate that text was just a necessary evil that visitors would avoid reading if they could.

Then my boss hired a professional copywriter.

I quickly realised that I had been working under the assumption that people don't like to read because it's a boring chore. It turns out that what people don't like to read is bad, irrelevant, unintelligible or over-the-top copy (i.e. the kind I'd been producing). When copy is well-crafted it engages your visitors, keeping them interested in what you have to say. While pictures may encapsulate a thousands words, the right three words may be easier for your customers to memorise, remember, and pass on.

These are my rules of thumb for having good copy:

- **Get a professional to help you**

If you're a professional designer, think of how much better you are at using Illustrator or Photoshop or a pencil to produce a design than your boss is using PowerPoint and clip art. That is exactly how much better than you a professional copywriter or technical author is at producing the right words.

- **Do not use dummy text**

The function of web and landing pages designed for conversion is to convey a strong persuasive message. If you're producing mock-ups, or worse prototypes, using dummy text you're starting the wrong way round. Instead start with the message and design around that.

- **Avoid superlatives**

Have you ever seen a first-timer using Photoshop? They'll go nuts with the effects producing headings and buttons rife with drop shadows, bevels, outer glows and, most probably, lens flare or two. The equivalent in writing copy is using words like 'most', 'amazing', 'outstanding' and similar superlatives that aren't descriptive and don't add value.

- **Don't be funny.**

Seriously. Successfully pulling off joke headlines and satirical commentary is really hard and when done amateurishly will at best result in an eye-roll.

- **Use page tables.**

Page tables are a useful technique for defining and remembering the messages (and their priority) your design is supposed to be communicating. The last section of the book contains a link to page table resources.

And finally, the litmus test: would you be willing to read out what you've written out loud? Do not ship, publish or print any design that contains a message you would be embarrassed of reading out loud to your partner/spouse/boss/best friend.

Rule 5: Aim for substance over significance

As you get ready to run a series of optimisation experiments and measure the effects on your conversion rate you will inevitably come up against the issue of significance. More specifically, statistical significance. A result is said to be 'statistically significant' if there is a high confidence that the result was not a consequence of chance.

For instance, let's say you have a landing page that currently gives a conversion rate of 10%. You design an alternative version and run a split test. After a week, the original version is still converting at 10% while your new version is converting at 11%.

Do you stop the test and replace the old with the new, do you stop the test and keep the old version or do you keep running the test for another week?

The answer to a large extent depends on other bits of data that I haven't given you here such as the volume of traffic going to both those pages, the volume of sales and whether that original 10% was a very consistent rate. There are a ton of online tools and books that will help you figure that out. By plugging in the data into a formula you can get a confidence level – usually expressed as a number – which if greater than a particular threshold means that the result was not a consequence of chance. If the confidence is high then you can roll out your better performing design.

But that's only half the story. Statistical significance should not be your endgame; substance should be. By 'substance' here I mean 'importance' because significance in statistical terms does not necessarily imply importance in real terms. Statistical significance may be a prerequisite for determining importance but its often not the only thing that matters.

These are my rules of thumb for determining whether the optimisation work you're doing is substantive:

- **Do you know why it happened?**

You know *what* happened: you observed a 50% increase in your conversion rate because you changed a headline or a layout. But do you know *why*? Do you understand what in your new design or message caused the change in customer behaviour: Is the message more persuasive? Is the design more usable? Is the pricing explained in a more attractive fashion?

When a test concludes with a significant result it is important to follow up with customers to see what really caused them to convert (rather than just going with your own assumptions). Building up this knowledge is what will

help you build your own set of best practices that are specific to your audience or business.

- **Is your result repeatable?**

Let's say your 50% increase in conversion rate came from optimising a landing page. That means that for every 10 units the old page sold the new page now sells 15 and the numbers show that result to be significant. However, can you repeat that success on your other landing pages? Is the change you made one that you can roll out to other pages and achieve the same result thereby multiplying the effect of the test?

Even a small company may have hundreds of landing pages each tied to different products and advertising campaigns. If you need to run dedicated and detailed tests to optimise each one individually your results, while significant, may be less substantive. A more important result may be one where you can apply what you learnt from tests on one landing or product page to *all* your pages thereby getting the benefit of that 50% increase across your entire company's sales efforts.

- **Is your result consistent over time?**

You may have observed that 50% increase in conversion after running a split test for a week or two and as we said the numbers show it to be statistically significant. As a result you roll out your new design as the new default. But for it to be a truly substantive result that increase would need to be consistently observable long after your test ends.

If a month into the new design your conversion rate baseline goes back down to the previous rate then your result was not substantive because the effects of your work are not long-term.

Calls to Action

So how do you use all of this stuff that I've been talking about here?

As with most of these sorts of books not everything will apply to your specific circumstances and you most certainly can't do all of it at once – even if you agreed with everything I suggested! Some things you probably knew already, like the importance of talking to customers. Others you may need to dig much deeper into, like crafting great words or understanding statistics better. And a few things you may be able to make progress on the moment you put down this book, like making sure you have the right goals in place for your conversion optimisation work.

Wherever you fall on that spectrum what I hope you go away with are two key principles:

The first is that designing to sell – let's use the S-word – can be fun and they don't have to be something you do reluctantly or warily. You can apply many principles coming from good user experience design to make the customer's experience of your online sales process a rewarding one for them and one that you can take pride in.

The second is that you need to build up your own set of best practices for the context that you're working in. There are very few silver bullets that will magically improve your conversion rates without a good understanding of the sorts of customers you're dealing with. Most of the ideas and guidelines that I've presented you with here aren't of the "put a button on the left-hand side" kind. Instead they're ones which will enable you to build an arsenal of knowledge and insights with which you can optimise conversions in a repeatable and scalable manner.