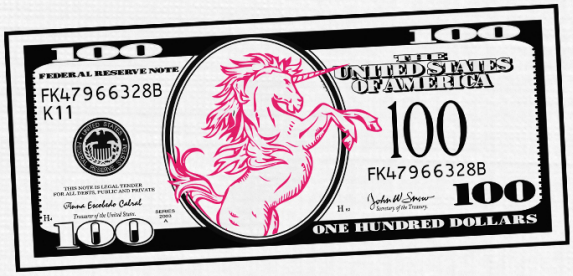
JACK KINSELLA







Entrepre NERD



The Busy Person's Primer on
SEO, Paid Advertising
AND Conversion Optimisation

www.dbooks.org

Copyright Warning

This document is secured with visible and invisible watermarks that connect this copy with the identity of the person who bought it. Legal action will be pursued if this book is uploaded online or made available on BitTorrent.

The author has repeatedly taken action against infringers in the past.

All rights reserved. No part of this publication may be reproduced, distributed, or transmitted in any form or by any means, including photocopying, recording, or other electronic or mechanical methods, without the prior written permission of the publisher, except in the case of brief quotations embodied in critical reviews and certain other noncommercial uses permitted by copyright law. For permission requests, contact the publisher.

Copyright © 2016 by Jack Kinsella.

Acknowledgements

The content of this book has been reviewed by a few enthusiastic readers who donated their time and thoughts. I'd like to especially thank Aditya Balachander, Thomas Carney, George Dallas, Frederik Eichelbaum, Fabian Harmik Stelzer, Heinrich Held, Stefanie Jakobi, Sascha Konietzke, Daan Loening, Adam Niedzielski, José Quesada, Julio Santos, and Benjamin Wilson.

Along the way I had help from various editors: Joseph Doherty-Bailey, Matt Rott, and Kai Wu.

The book was produced and typeset automatically using the excellent open-source Pandoc, Calibre, and LaTeX libraries.

Errata Reporting

Rare is the book that goes to press without errors. It is only with the help of the first batch of readers that a manuscript reaches

its most polished form. For that reason, I'd be greatly in your debt if you would report any typos, formatting issues, or factual inaccuracies you encounter along the way.

http://www.jackkinsella.ie/customer_support_requests/new

www.dbooks.org

Contents

Copyright Warning 2
Acknowledgements 2
Errata Reporting 2

1 Introduction **12**

PART A: SEARCH ENGINE OPTIMISATION

2 Intro to SEO **14**

3 SEO: First Steps **16**
Research and Choose Keywords 16
Bonus Points: Piggyback off your competitors 18
Buy Keyword-Based Domain Names 19
Choose Your Top-Level Domain Carefully 20
Avoid Keyword Cannibalisation 21
Concoct a Scalable Content Creation Strategy 22
i. Incentivise your users into generating public content 23
ii. Automatically generate content from algorithms . 25
iii. Hire content generators 26

4 SEO on Each Individual Page **28**
Ensure You Can Satisfy a Visitor's Searching Intention . 28
Research What Currently Ranks 29
Place Your Keyword in the URLs 29
Place Your Keywords in the Title Tag 31
Include Your Keyword Early in the HTML 32
Carefully Craft a Unique Meta Description 33
Ensure Your Website Makes Sense to Googlebot 34
Include Rich Snippets and Structured Data 35
Every Page Should Include a Picture 38
Write Longer Pages 39

5 SEO at the Whole-Website Level **40**
Generate a Sitemap 40

Link Internally 40
Use HTTPS Encryption 42

6 SEO by Mitigating Duplicate Content Issues 44
Redirect Instead of Responding 46
Use Canonicalization 47
Write a Robots.txt or Meta Robots Tag 48

7 SEO in Image Search 52
Output Public-Facing Images Through Your Software . . 52
Add Alt Text to Your Images 53
Choose Informative Filenames for Your Images 53
Write Captions to Accompany Your Image 53
Generate and Submit an Image Sitemap 54

8 SEO Through Backlinks 55
Include Contributions from People with Big Audiences . 56
Coin Concepts 57
Leave Your Link Where You Can 57
Email Outreach 58
Pitch Your Content To Writers of Round-Up Articles . . . 58
Piggyback off High-Ranking Websites 59
Guest Post on a Large Platform 59
(Maybe) Buy Links 60

9 SEO Vis-à-Vis Social Media 61
Include Open Graph Tags 61
Include Twitter Cards 64
Promote Your Content on Social Media More Than Once 64

10 SEO and Internationalisation 65
Distinguish Language from Territory in Architecture . . 65
Choose a Fitting Structure for Internationalisation 66
Communicate Your Structure to Google 69
The Hreflang attribute 69
Viewing your internationalised Google search results 72

PART B: COPYWRITING

www.dbooks.org

11 Copywriting **73**
Emphasise Benefits Over Features 74
Convincingly Demonstrate Credibility 78
Give Reasons to Act Now as Opposed to Later 83
Prefer Concrete, Visual, Specific Copy 86
Ensure Scannability, Readability, and Responsiveness . . 89
Signal with Indirect or Unworded Messages 90
Explicitly Point Out Qualities, Lest They Go Unnoticed . 91
Compare Yourself to the Competition 94
Anticipate and Counter Objections 96
Bonus Points: Mention a negative, then a positive. . 97
Associate Yourself with Desirable Connotations 97
Repeat Yourself 100
Don't Talk Yourself out of Markets 102
Keep the Scent 103
Arouse Curiosity 104
Play on Negativity 107
Anchor Numbers 108
Give Detailed Information 109
Depict Yourself as Someone Similar to Your Readers . . . 109

PART C: ANALYTICS

12 Google Analytics **111**
Why Bother with Analytics? 111
Workflow Tips 113
The sanity-preserving Analytics workflow: Form plain English sentences backed by hard data . 113
Navigate Analytics by searching 114
Connect to Other Google Products 114
AdWords 114
Search Console 116
Location of Snippet in HTML 116
Cross-Browser and Cross-Device Tracking (User ID) . . . 117
Subdomain Blues 119
Conflation of traffic with the same relative paths on different subdomains 119
Session data resetting across subdomains 120
Cleanup Noise, Chaff, and Distortion 121

Lowercase Request URL 121
Lowercase all UTM parameters (URL Builder parameters) 122
Remove sorting/pagination/noisy parameters 122
Exclude bots 125
Exclude internal IP address from your live tracker . 125
Set up a separate web property in Analytics for staging/testing 126
Exclude unwanted referrers 126
Analytics and Social Media 127
Social interactions 127
Custom Dimensions 128
Illuminating Analytics Blind Spots 129
Virtual pageviews 129
Track important occurrences that don't correspond with page views 131
Collect Your On-Site Search Queries 133
Measuring Success 134
Setup goals 134
Track ecommerce transactions 135
Tag URLs 135
Precautions 140
Be Disciplined in Making Annotations 141
Send Alerts 142
Analytics Blocking 142

PART D: CONVERSION OPTIMISATION

13 Conversion Optimisation 144
Its Advantages 144
Underrated 145
On-site vs Off-site 146
Beware of Shallowly Applying the Lessons of Others . . . 147
Areas Worthy of Experimentation 148
A. Conversion funnel length 148
B. Funnel order 152
C. Number of clickables per page 154
D. On-page copy 156
E. Images 157

www.dbooks.org

F. Pricing 159
G. Timing 162
H. Salience (includes graphic design) 162
I. Choice architecture 164
J. Available features 165

PART E: STATISTICAL SIGNIFICANCE

14 Statistical (In)Significance 167
The Theory 167
The Practice 173
Know what counts as a sufficiently "large number" in the "Law of Large Numbers" 173
The difference between ultimate and intermediary goals. 177
Understand the scope and business value of your results. 178
Run your tests in parallel, not consecutively. 181
Do not peek at your test results early. 182
Use a statistical significance calculator. 183

PART F: EMAIL MARKETING

15 Email Marketing 184
The Law 185
Best Practices 186
Classic Promotional Emails 187
Sale 187
Cart abandoners 187
Recent churners 188

PART G: PAID ADVERTISING

16 Introduction to Paid Advertising 189
Advertisement Is an Annoyance 189
No One Clicks on Paid Adverts 191
Doesn't Everyone Have Ad Blockers, Anyway? 192
Paying for Attention Is Dirty 193
It's Too Expensive 194
Why Pay for Traffic That Would Have Visited Anyway? . 195

1. The cost isn't that high. 196
2. Countermeasure to competitor advertising. 196
3. Control over advertising messages. 196
4. Social signalling of trust/credibility. 197

17 Paid Advertising: First Steps 198
Building an "Advertising-Ready" Web Application 198
1. Create appropriate landing page constellations . . 198
2. Calculate profitability with code 199
3. Write automated tests for conversion pixels (etc.) 199
4. Only advertise to permanent URLs 200
5. Place conversion pixels high up in the HTML of the page 200
Analytics and Advertising: The Eyes to See With 200
Conversion tracking 203
Pay Per Click? Per Impression? Per Conversion? 208

18 Paid Advertising: Staying Organised 211
Structuring and Labelling Your Accounts Sensibly 211
Campaigns 211
Ad groups 213
Ads 215
Don't Delete Data 216
Change History 216
Schedule Review in Future 217

19 Paid Advertising: Targeting 219
General Principles 221
Off-target targeting 221
Dressing for the occasion 222
The AND gotcha—avoiding unnecessary targeting . 222
Oblique targeting 224
Negative targeting 225
The edge 227
Synonym bombing 228
Undersubscribed targeting groups 229
The sum is sometimes more than the parts 229
Research advertising ease before committing 230

www.dbooks.org

A Tour of Targeting Options 231
Geographical location 231
Time 232
Gender 233
Age 234
Device/computer/browsing technology 235
Interests (/ Likes) 235
Keywords 238
Your fans 241
Other people's fans 243
Uploaded lists of customers/leads 243
Behaviour 246
YouTube 247
TV 248
Major world event 248
Relationships and family 249
Users of your own app 250
Users of other people's apps 251
Visitors of your own website (remarketing) 252
Visitors of other websites 256
Language 258
Education 260
Workplace and job targeting 260

20 Paid Advertising: Ad Creatives 262
Image Choice 262
Headline 266
Body 266

21 Paid Advertising: Strategic Principles 268
Topical Advertising vs Interruption Advertising 268
Rationed Views 269
1. Possibility of gas pedals 270
2. Targeting migration into the underlying group . . 271
3. The need for anti-advertising-blindness measures 271
4. Frequency capping 272
The Progressively Rising Cost of Additional Leads 275
Switch On All Bells and Whistles 278

Unintentionally Competing with Yourself 279
Staff Working for Advertising Platforms Are Experts in Platform Features, Not in Your Profitability 280

22 Paid Advertising: Debugging and Optimisation 282
What Advertising Platforms Want 282
Stay Empirical 283
How Much Should I Start Off Bidding With? 283
What's the *Most* I Should Ever Bid? 285
When Metrics Mislead You 285
1. Misleading CTR (click-through rate) 285
2. Misleading CPM (cost per 1,000 impressions) . . . 286
3. Misleading CPC (cost per click) 287
4. Misleading CPA (cost per acquisition) 288
5. Misleading ROAS (return on advertising spend) . . 289
Overcoming Low Reach 290
i. Targeting woes 290
ii. Overaggressive filters 291
iii. Tough competition 291
iv. Low Quality Score 292
v. Tightly rationed audience 293
Overcoming Low CTR 293
i. Fine-tuning targeting 294
ii. Completely different targeting 295
iii. Form of advertisements 295
iv. Content of advert creatives 296
v. Match between targeting and advert content 296
Overcoming Low Quality Score (Google-Specific) 297
Overcoming Low Relevance Score (Facebook-Specific) . 298
Overcoming High CPC 299
Overcoming High CPA 300
The End 301
Review Request 301

www.dbooks.org

1 | Introduction

In a fair world, meritocracy would rule. Being good at what you do would be all that's necessary to get fans, sales, custom, and acclaim.

But we do not live in such a world; our reality is a crowded one, one where only those who shout the loudest get noticed.

It is disconcerting to consider our own insignificance, but it is only by taking this perspective that we grasp the difficultly of entering the public's awareness.

Given vast access to the machinery of the web and digital publishing, each of us is competing for attention with a billion others. Each of us has but one painting sitting in a gallery housing a thousand others, stationed in a city hosting a thousand such galleries, situated in a world with a thousand such cities. Leaving our painting's discovery to fate alone, our chances of being noticed are little better than those of winning the lotto.

This state of affairs is hardly a new one. Charles Ives, considered by some to be the greatest American composer, ceased writing music 18 years before his death. For the remainder of his life, he dedicated himself to promoting and distributing his music. Without this gargantuan effort, it is doubtful that his work would have the acclaim it enjoys today.

The ultimate goal of all marketing is to get ever-increasing, self-perpetuating word of mouth (or, to be more modern about it, "recommendation of Reddit" or "tweet of Twitter"). But this is a case of the rich getting richer, of attention begetting attention, of the virtuous circle in full swing. Because word of mouth only kicks in after the snowball is already hurtling down the slope, a more apt question is how to get the thing to budge in the first place.

The approach outlined in this book is one possible answer to this question. In a nutshell, the strategy is as follows: You will take out online advertisements and invest in SEO so as to inform relevant members of the public that you exist. Thanks to the potency of these measures, a steady stream of promising leads will arrive to your website. Here, you will apply the psychological tactics of conversion optimisation and copywriting to convince these leads to place orders, download your music, or whatever it is you want them to do. Following their conversion, you will send these past customers automated emails and display remarketing ads so as to keep them coming back to you again and again. Meanwhile, throughout every stage of this process, you will study reports in Google Analytics to discover which particular marketing measures brought you the highest profits. So as not to be fooled by randomness, you will apply statistical significance testing, thereby ensuring that any confidence you have in your marketing conclusions is justified.

I want to highlight that not all of your marketing forays will bear fruit. Indeed, most will fail. The reality is that you will need to try out lots of different angles until you find the one that sticks.

And with that, let's dive in.

www.dbooks.org

2 | Intro to SEO

Search engine optimisation (SEO) is about more than just ranking #1 in Google. Obviously that's a big part of it all, but limiting our vision to just this component would ignore other important aspects of the domain.

For one, topping the ranks in Google won't do you one iota of good unless your page's search result snippets compel people to click through—and whether or not they do this depends on how enticing you've made your snippet, a factor you control.

And there's not just Google Search anymore, but also Google Video, Google Maps, and Google WhateverComesNext. These are separate "search races," and admission to Google Search basic doesn't automatically grant admission to any of the others.

Not to mention all the search engines beyond the Googleplex—and by this I don't just mean competing engines like Bing or Yahoo, but also niche-specific engines like Amazon's product search, Yelp's restaurant search, TripAdvisor's hotel search, or any search on a major third-party website that could potentially send you traffic. From this perspective, SEO is partly about swapping the single-entry ticket for the season pass.

There's a big bonus for carefully SEO-ed pages: Many optimisations, particularly those affecting the on-page ranking factors, also assist in raising the so-called "quality score" in Google Adwords too, thereby lowering the cost of online advertising. That means SEO has a double payoff.

So, I believe this section contains the highest impact information in this book, simply because SEO can be mind-blowingly cost-effective compared to any other marketing technique.

But cost-effective does not mean free. Most webmasters fail at

SEO because they fail to invest serious resources into it. By "resources", I don't just mean money—I'm also referring to the brain-cycles of your best staff. SEO is like buying a house, whereas paid advertising is like renting one. Buying involves a high upfront cost, but afterwards you get to live rent-free. Renting, by contrast, has low upfront costs, but leads to a larger cumulative bill over the course of a lifetime.

Doing SEO *right* requires you to spend the resources—and if you aren't prepared to do this, you might as well keep renting traffic for the rest of your life.

Disclaimer: No one outside of Google's search team can say with certainty what will or won't work for SEO, and so what follows is a combination of officially sanctioned "best practices", industry whispers, and conjecture based on observation and common sense.

www.dbooks.org

3 | SEO: First Steps

Research and Choose Keywords

The core of SEO is keyword analysis. Even before choosing a business name or buying a domain, the savvy SEO expert first researches the language and vocabulary used by the biggest proportion of their customers. Do more people search for "radio-controlled airplanes", "rc airplanes", or "flying airplane kits"? In order to answer this question, the SEO expert seeks out publicly available data, principally in Google's Keyword Planner, so as to ground their marketing decisions in quantitative data. By knowing what keywords people type into search platforms, the marketer can build their strategy upon a solid foundation.

The marketer chooses a set of keywords with the intention of ranking highly in Google whenever these keywords are searched for. But because ranking highly is unlikely to happen by accident, the marketer will need to bake these keywords into their website's information architecture and design, for example by including the keywords in their URLs, page titles, and internal links. In addition to this *internal* (or *on-site)* SEO, the marketer will ask *external* (or *off-site)* websites to link to them (*backlinks*) by including the keyword as part of the external link's anchor text. After all this is done, these keywords will get one more use: They will form the basis for any paid advertising that employs keyword targeting.

Novice internet marketers often believe that the best keywords are the ones with the most monthly searches, as estimated in Google's Keyword Planner. But this isn't always true. When someone googles something generic like "financing", there may be all sorts of different intentions behind their search. Is it a young family looking into financing a home? Is it a car dealership looking to partner with a bank to provide their customers with financing? Or is it a student researching the topic for an economics essay?

As you can see, the keyword "financing" means a lot of things to a lot of people. If a bank that sells mortgages were to attempt to rank for this keyword, then they would not only face heavy competition from similarly minded car financiers, but also receive a lot of irrelevant traffic that bounces instantly, an occurrence which (as I explain later) also hurts SEO efforts.

The marketer should think in terms of the commercial intent behind the search, as relevant to their business. That means that the bank above would only choose keywords specifically relating to home loans, as opposed to broader financing.

Their exact choice of keywords will hinge on how they trade off strategic factors. If they are just starting up their home loan business, there's no realistic chance that they will even rank on the first *page* of search results for something as competitive as "home loan financing", not to mention ranking #1 on that initial page. Instead, they would be better off choosing their battles wisely and investing their efforts into a more specialized keyword for which they have a realistic chance of ranking.

In strategically weighing the pros and cons of each keyword, the marketer will pay attention to the "Suggested Bid" column in Google's Keyword Planner. This figure represents what Google thinks it would cost to advertise for this keyword on Google Adwords. Assuming that competitors are rational actors, the more expensive keywords are likely to generate more valuable traffic. So why not try to rank organically for these hot leads? (Proviso: High advertising fees mean that the keyword is valuable to *someone,* but not necessarily to someone selling whatever it is you offer.)

Yet another element of strategically choosing keywords involves sizing up the competition. The marketer will google each of their candidate keywords to see what kinds of pages currently rank highly. Because content is ultimately king in SEO, the marketer needs to be confident that they can produce something that blows the current ranking champions out of the water—not just in terms of quality of writing, but also in comprehensiveness, as

www.dbooks.org

well as impressiveness of accompanying illustrations and design.

Those of you with international ambitions, be warned: The optimal keyword might vary by country, meaning that not only must you be prepared to consult the data again when entering a new territory, but you also should have designed your application to support keyword-switching depending on the region. In my business that sells study notes, I erred by hard-coding the keyword "{X} notes" into practically every layer of my application, only to find that once I started operating in the USA, "{X} outlines" performed better. I found myself in the midst of a messy refactor due to that lack of foresight.

So far we've focused on *direct-route* keywords—descriptive keywords that map onto whatever the business sells (e.g., a piano seller targeting "buy pianos" or "piano showrooms"). But sometimes direct-route keywords are not a viable strategy—say if the competition is too steep, or if the search volume is too low. In this scenario, it's up to the marketer to figure out *indirect* keywords, which map onto lateral activities and interests of the business's customers (e.g., now the piano seller targets "piano practice regimes" or "local piano teacher listings"). This approach relies on the idea that people interested in X are also often interested in Y, and the marketer can learn about these correlations by observing what their audience blogs about, reads about in their spare time, shares on Facebook/Twitter, talks about on Reddit, and so on.

Bonus Points: Piggyback off your competitors

I know of a successful "how-to" publishing company that optimised its page for the keyword "{X} NOT For Dummies", with the intention of siphoning off traffic that originally searched for books by their established competitor, as would be evidenced by searches such as "X For Dummies". This upstart now ranks #5 on searches for their much larger and better-funded competitor—not bad.

Buy Keyword-Based Domain Names

Keywords contained in the domain name count towards SEO, especially when the domain name is an exact match for the query (i.e., where there are no other words present).[1] So, if you're selling German bread, "tom-d-germanbread.co.uk" is better than a totally unrelated domain name like "tom-d-services.co.uk", and "germanbread.co.uk" is better than "tom-d-germanbread.co.uk".

There is a downside to keyword-based domain names: They are difficult to brand, and this has negative effects on many aspects of marketing. That said, there may be a way out of this problem: Buy both a branded domain name (i.e., one that is easy to spell, easy to remember, etc.) and lots of keyword-based domains. Next, configure each of these keyword-based domains to run on the same server, albeit with a different "skin". This lets you present various keyword-specific "mini-sites" to the consumer, each of which will rank well for that one keyword. (It is advisable for these websites *not* to link with one another though, lest you give Google the impression that you are running a link farm). For an example of how this would work, look to Patrick McKenzie, creator of BingoCardCreator.com (a web application that helps teachers print Bingo cards for their students). A few years ago he bought a sizable collection of exact match domains for his highest-performing keywords. For example, to help himself rank for "Halloween Bingo Cards", he bought "http://www.halloweenbingocards.net/" and created a special HalloweenBingoCards mini-site. For his trouble (and the 10 bucks it cost him to pay for the domain name), he scored himself 27,000 hits that Halloween, triggering 15 sales at $35 a pop. This works out as a total of $535 in recurring sales for the $10 he pays yearly for the domain name. Now imagine doing this en masse with *hundreds* of domain names in a business that isn't quite so niche.

[1] https://moz.com/search-ranking-factors/correlations#4

www.dbooks.org

Choose Your Top-Level Domain Carefully

Top-level domains—the part after the last dot in the domain name (e.g., .com, .org, .co.uk, .de)—are categorised as either bound to a particular country (.co.uk, .de) or unattached (.com, .org). If your initial domain was country-specific but you later wish to internationalise your business, you're going to end up with headaches.

Case in point: I started my business with the domain-name "oxbridgenotes.co.uk". At that time, my business only served British customers and I had limited business ambitions, so I thought that a .co.uk would be A-OK. But with time my ambitions outgrew my UK-bound domain, and, to my dismay, I discovered that any suitably global URL structure that I could create (e.g., subdomains like "fr.oxbridgenotes.co.uk" or "de.oxbridgenotes.co.uk") would not only look weird to customers in these new regions, but also fail at SEO because search engines predominantly display .co.uk content to UK searchers. This remains true even if contrary information is given in the subdomain or folder information (i.e., "oxbridgenotes.co.uk/fr" still displays in the UK, not France). All in all, the top-level domain dominates the subdomain in terms of SEO signalling to both search engines and humans, and there isn't a whole lot one can do about it.

In order to move beyond the UK, I had to buy a new domain, this time a generic, non-country-specific one ("oxbridgenotes.com"). From this base, I could add country-specific subdomains or folders as I wished, and Google would respect this structure in their rankings.

But there was a big downside to changing domains so late in my product cycle. My .co.uk domain had been live for years and attracted good organic traffic. Naturally, I didn't want to shut it down. In the end, I decided to maintain both the old domain ("oxbridgenotes.co.uk") and the new one ("oxbridgenotes.com"), running the pair on the one server, configured to grudgingly gel together via an annoyingly complicated and bug-prone routing

setup.

Aside from programming complication, this splitting of my website into two domains was far from ideal for another reason: I ended up spreading my "SEO juice" across two different domain names—the old .co.uk one, and the .com which was virgin and therefore weak in search rankings. A new product released on the .com wouldn't rank at all, whereas the same thing on the .co.uk would do very well. Had I started off with a single domain name from the very start, then I wouldn't have had this two-tier SEO and brand problem; everything released onto the website would have ranked respectably.

(I am aware that there are facilities to officially move domain through Google Search Console, but I was and still am too risk-averse to toy with this particularly treacherous type of fire.)

In conclusion, I would advise my former self to be a little more optimistic about eventual international expansion and start the website on a top-level domain that isn't tied to a particular country.

Avoid Keyword Cannibalisation

Google views the web as individual pages, not as entire domains. So, contrary to popular myth, there is no such thing as ranking your *domain* for a keyword. In fact, what's really ranking there is your homepage, which often gets the most internal and external links.

Instead, you ought to think of your website as a collection of individual pages that compete with each other and every other web page out there to rank for some particular keyword or keyphrase. This competition between pages on your own domain leads to a problem known as "keyword cannibalisation".

Imagine a webmaster slapping the keyword "wordpress templates" on 200 different product pages of their digital downloads website. Then no single product page is obviously pre-eminent

www.dbooks.org

to a third party wishing to direct their customers to this source of wordpress templates. As such, whenever a third party links to our webmaster, the link could point to any one of these 200 pages instead of concentrating on the most relevant one. The end result is that none of the pages on the digital downloads domain reach the strength needed to appear on the first page of Google results.

Furthermore, this sort of "tactic" is damaging to the digital download seller's overall conversion rates. Some product pages might be better at converting visitors than others, yet the better and the average pages show up side-by-side in Google search results. As such, a random visitor is as likely to land on a page that's poor at converting as a page that's a conversion winner.

The webmaster would be better advised to choose one unique compound keyword for each page, such as "wordpress templates for pubs", "bootstrap wordpress templates", or "easy-to-use wordpress templates", and focus all their on-page SEO signals—the URL, the title tag, etc.—on this narrower goal.

What about the original catch-all parent keyword, "wordpress templates"? The webmaster should invest in creating a stellar web page for this compound keyword, something that people find very valuable and therefore link-worthy. Then, they should ensure that every page focused on narrower compound keywords should link back to this parent page, thereby attributing some SEO juice via internal linking.

Now our webmaster has 199 internal links to their "wordpress templates" page, and an obvious catch-all place for third parties to link to—a useful result supporting general and specific SEO!

Concoct a Scalable Content Creation Strategy

Roughly speaking, the more unique and useful content your website has, the better you'll do. A full-fledged SEO campaign therefore includes an intelligent and cost-effective way to generate this precious basic ingredient of ranking. This idea is eloquently captured in a buzzword coined by Patrick McKenzie: scalable con-

tent creation.[2]

Your specific choice of scalable content creation strategy will affect the eventual form your business takes, so it's important for you to choose your course before writing your first line of code. The key piece of homework is answering the question, "How do I create large amounts of content *valuable to customers* at a low cost of time, effort, and/or money?"

Notice my emphasis on "valuable to customers"—too many websites forget this and pump out any old drivel for content. This is unwise though, because Google can detect dissatisfied visitors by the fact that they press the Back button, by their time spent on your page, and/or by their repeating their previous search. If this happens often, Google interprets this as a strong signal that your web page sucks, and they respond by banishing you to the lowest ranks of the search ladder.

One option to create content is to write it yourself. This can be a good idea to start with, but it doesn't *scale*. Were you to spend your entire working day writing, you might only be able to produce at a rate of five pages per day, while not investing time in the rest of your business. Not ideal. So what are your other options?

i. Incentivise your users into generating public content

Whereas you are but one person, your users may number in the thousands or even in the millions. By giving them some advantage or value (which may be as simple as a place to express themselves and be heard), you can incentivise them to create content on your platform—content that will be picked up by search engines and employed to attract even more users, who will themselves produce content, continuing and amplifying your reach ad infinitum.

There are all sorts of user-generated content creation systems out there in the wild, and I'd like to share some of my favourites:

[2] http://www.kalzumeus.com/2010/01/24/startup-seo/

www.dbooks.org

1. Github.com offers a free code-hosting service for anyone with open-source code projects. The public activity of these legions of programmers leaves behind millions of pages of content, composed of code project readmes (which Github automatically converts into HTML), issue-tracking discussions, code review comments, and so on. Github gets free content, and the programmers get a free place to host their open-source code while showing off their coding chops to their peers. Everyone wins.

2. People love giving opinions, and by playing to this innate human drive, websites like Amazon generate an unending stream of content by asking anyone (not just their customers) to review the products they sell. And considering that Amazon sells pretty much everything, this amounts to a whole boatload of review content. Another great source of user-generated content are the commenting sections underneath online newspaper articles. Sensitive topics—like Palestine or migration—often compel thousands of readers to contribute their opinions. To see this content-creation strategy done at its finest, check out the Guardian's website. Sometimes every word of official editorial content yields over a hundred words in public commentary.

3. In my primary business, Oxbridge Notes, the average author we publish on our platform uploads about 100 pages of notes. These well-organized packages of pages constitute an eloquent lode of user-generated content that has never been on the internet before. The customers who buy these packages like to view free samples before buying, so we wrote an algorithm that turns the first 20% of each uploaded package into such a sample. This gives us a yield of about 20 pages of unique content per author. To capitalise on this content, we create landing pages for each of these free samples, complete with a PDF and an HTML transcription of the text. This strategy is completely automated, yet accounts for 94% of our organic traffic. The average time spent on each of these generated pages is 1.5 minutes, showing that customers find the text there to be valuable.

4. YouTube and SoundCloud tap into people's desire to express

their creativity and share the results with the world. By providing platforms for hosting video and music, both websites have drawn huge, global audiences. And if that weren't already enough, these platforms also offer comment sections where admirers and haters duke it out regularly for days and sometimes even years, leaving behind even more indexable content.

5. Stack Overflow, a popular question and answer forum for programmers, envelops both its questioners and its question-answerers with addictive gamification, consisting of reputation points, badges, and specially unlocked website privileges. But that's not all—not by a long shot. Stack Overflow rewards their question-answerers *emotionally* through granting them "proof" that they are right/helpful, via the mechanism of a questioner "accepting" one answer above all other answers from other contributors. To top it all off, Stack Overflow has convinced many people that these reputation points count towards an employer's hiring decision, enabling Stack Overflow to turbocharge their content generation with the fuel of career ambition (while boosting their business model: job ads). Indeed, some companies employing programmers have started weighing Stack Overflow reputation points in determining a prospective employee's ability, so this approach has become a self-fulfilling prophesy.

6. Ojbc.io, an online magazine for iPhone programmers, started by building a slick-looking web platform and writing their first issue with in-house authors. They then polished their content with the help of a professional editor and the addition of a pretty magazine cover. This distinguished them from your average blog, and they were able to leverage this edge to convince thought leaders to contribute articles to their subsequent magazine issues. In exchange for their initial efforts with that first issue, they managed to convince brilliant people to contribute world-class content to the subsequent 23.

ii. Automatically generate content from algorithms

Automatically generating content (e.g., by combining database information, scraping, and clever algorithms) often only results

www.dbooks.org

in reprehensible, unreadable spam-level output. That said, in rare cases a clever webmaster has the skills and taste to pull off automatic/algorithmic content generation, and they may benefit magnificently for their efforts.

One such example is Versus.com, which generates English language sentences from data points about comparable electronic devices (e.g., phones or video game consoles). These data points, presumably gathered from manufacturer websites or massive retailers, are about things like the relative battery life of devices, their processing power, their screen size, or their backwards compatibility. Relying on this strategy of turning API-consumable data points into HTML sentences useful to humans, they have acquired monthly traffic in the millions. And their content, despite having been generated by an algorithm, is sufficiently compelling to ensure long and often heated discussions by real humans in the comment section appended to each page.

iii. Hire content generators

Patrick McKenzie hired a school teacher to write bingo cards on given topics for his website Bingo Card Creator. This school teacher produced 30 sets of cards every month through a custom CMS that Patrick built into his web application. Once these bingo cards were uploaded, the Bingo Card Creator website algorithmically turned them into search-engine optimised web pages.

The key to this strategy was the transformation of content creation into a repeatable task that could be performed by a non-technical collaborator and repeated thousands of times. After a while, this built up a formidable bank of content, all without Patrick having to contribute any more of his own time.

A similar trend in the ecommerce world involves hooking freelance writers into a website's CMS so as to write unique product descriptions, say for the cameras or the phones being sold. The motivation here is that ecommerce websites relying only on the manufacturer's product description often have product pages with insufficient word counts for Google to bother index-

ing them. And even if Google *did* index these pages, there would probably be insufficient unique content to outrank the gazillion other retailers selling the exact same product with the exact same description texts.

www.dbooks.org

4 | SEO on Each Individual Page

Ensure You Can Satisfy a Visitor's Searching Intention

One of the major factors that search engines use to determine how highly you'll rank in Google is whether or not your web page manages to satisfy a visitor's search query. Imagine someone searching for a "private math tutor in Manchester". A website promising the same appears at rank #1. The googler clicks through and finds a skeleton page that is stuffed with keywords for "private math tutor in Manchester" yet doesn't have a single tutor available. The searcher is pissed off, hits the back button, and then clicks on another search result, this time the one at rank #2. This second website has a great selection of available tutors, so the visitor engages one of them and ends their search.

The next time this person carries out a search on Google, it will be for something unrelated (e.g., "how to build a giant catapult"). But all this while, Google has been listening with its ear to the ground. It knows that the searcher didn't find what they were looking for with the first website—otherwise, why would that person have hit back and searched for the same thing again? Google interprets this as dissatisfaction, and they respond by adjusting the future rankings of the two websites such that the one that satisfied the customer's need gets nudged upwards while the lamer alternative gets pushed downwards.

As a consequence of all this, strategic webmasters should hide or actively dissuade web surfers from visiting unsatisfactory pages. For example, the page for "private math tutor in manchester" should not have been created and indexed.

Research What Currently Ranks

There's nothing quite like imitation as a starting point in writing a web page intended to rank highly. Before writing the content for a given page, you are advised to study what currently ranks for that keyword. What format do these pages have—are they lists or essays or white papers? In what tone are they written? Who is the intended readership? What topics do they address?

Don't limit your research to Google Search either—check out what gets upvoted on social bookmarking websites, shared on social media websites[1], asked in question forums[2], or written about in books on Amazon.

Place Your Keyword in the URLs

There are a bunch of on-page SEO signals that you can influence to help you rank. One of the most prominent of these is the page's URL, and it is imperative that your URL contains your keyword.

Google's algorithms aside, there is also a human basis for this practice: Potential visitors encounter links to your web page while browsing other websites. They quickly speed-read the URL and make a snap judgement about whether it justifies opening up a new tab. If you include your keywords in your URL, you send a message-in-a-bottle to these potential visitors, a subtle communication that you might just have what they are looking for. The end effect is that you are more likely to pique their interest and cause them to click the URL leading to your website.

(Doesn't anchor text stop people from reading URLs? When available, yes. But anchor text is very often omitted (e.g., when someone copies and pastes a link into Twitter). This means that the raw URL will be shown to the potential visitor instead of the anchor text shell.)

There is a further reason to place keywords in your URLs. Google

[1]Buzzsumo is a great tool for researching social share data.
[2]WebpageFX's faqfox shows you what questions people ask in various Q&A forums.

www.dbooks.org

applies a special optimisation bonus within Google Search whenever your URL matches all or part of the searcher's query. They do this by bolding matching words in your URL, and this added salience increases the chances that potential visitors will notice your search entry amongst all the others, increasing the click-through rate from Google Search to your website.

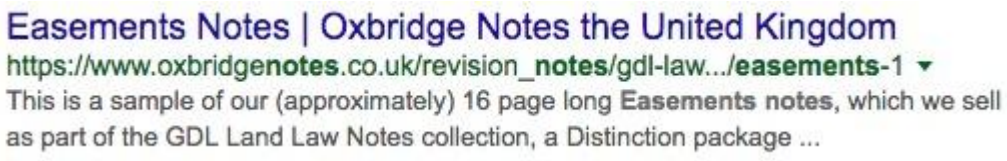


Figure 4.1 – *I searched for "easements notes". Within the search result entry, notice how "notes" and "easements" are both bolded.*

Try to avoid using machine gobbley-gook in your URLs. Imagine that the URL for the Easements Notes product pictured above was ".../search-db/GF38838383/easements#load-delay". Now, this web page runs the risk of matching for totally unrelated search queries, such as "search-db" or "load-delay". Moreover, all this ugliness in the URL is a waste of a good opportunity to put in human language that is more amenable to ranking, such as the student level of the notes ("GDL", a postgraduate legal degree) or the institution that the notes were created at (e.g., "University of Laws").

I'd like to give some advice to those in control of websites where at least some of the URLs are generated by administrators or regular users: In these circumstances, there is the danger of amateurly chosen URL names, URL deletions, and keyword cannibalisation.

- In response to amateurly chosen URL names, you might want to have a trained administrator touch up user-generated URLs before going live. Or, alternatively, you might build a software module that automatically inserts important keywords into URLs, ensuring that they stay relevant.
- Regarding deletions, you might consider some safeguards that prevent a user from deleting or renaming a URL that

once drew fantastic traffic. I personally have a system that keeps a history of old URL names and redirects requests for no-longer-existent URLs back to active URLs.
- With respect to keyword cannibalisation, the most basic (albeit underwhelming) precaution is a uniqueness check on URL names. A more comprehensive approach would involve the editorial touch of a human or a sophisticated NLP algorithm.

Place Your Keywords in the Title Tag

Google apportions heavy ranking weight to the title tag in your HTML, so one of the most basic points of SEO is to include your keyword(s) here.

As well as being used internally in Google's ranking algorithm, your title tag also appears as the title of your web page's search result snippet. For this reason, you should ensure that it communicates relevance to the human reader and entices them to click.

Due to the space constraints of a Google search result, your title tag is limited in length to approximately 55 characters. (An exact number cannot be provided because Google varies cut-off length depending on the device used to search.)

The title tag is typically used as the default text when someone shares your website on social media websites like Facebook or Twitter. By optimising it, you can better control and optimise how your page appears when shared virally.

The title tag also appears as the name of the tab/window in your browser. The marketing value of this little fact becomes apparent when we consider the visitor who gets distracted before buying and opens up Facebook in another tab, leaving your website aside for the time being. As they browse elsewhere, they will continue to see your title text in the corner of their eye. This serves as a reminder of your existence, or as a subtle to-do. There is perhaps a more subtle advantage to your peripherally present title. There is something called the "mere exposure" effect in psychol-

www.dbooks.org

ogy, whereby people develop preferences for things they are familiar with. This theory would predict that you can warm someone's feelings towards your brand through presenting it to them in their peripheral vision. You might think that this sounds like psychobabble, but Bernstein, R.F (1989) demonstrated in a meta-analysis of 208 experiments that not only is the mere-exposure effect robust and reliable, but its effect is strongest when unfamiliar stimuli are presented briefly.

If you want to do something to actively grab attention after being relegated to a tab, consider updating your title tag every so often (e.g., Facebook adds a (1) to the title whenever someone receives a notification). This change in tab title grabs people's attention, even from their peripheral vision, because the human eye is much more attuned to movement compared to stasis.

Google apparently weights the words at the start of your title more heavily than those at the end, so prioritise your keywords accordingly.[3] This means that websites which wish to include their brand name should slap it at the end of the title instead of at the start.

Include Your Keyword Early in the HTML

There is an old SEO practice known as "keyword stuffing" which consists of cramming your keyword as many times as humanly possible into a web page's body. This practice is now penalised by search engines, so it's no longer a good idea to write "cheap hats" 10 times within the page. That said, Google's algorithm *does* expect moderate mentioning of the keywords within the page's body—after all, if they penalised a web page about Germany for using the word "Germany", then the internet would become a giant cryptic crossword.

The latest search engine algorithms get around this problem by comparing a page's keyword density with that which one would expect in such a web page, meaning that artificial and unnatural

[3]https://moz.com/learn/seo/title-tag

keyword stuffing will stand out and be penalised. Playing it safe, then, your best bet is to write a web page that earnestly attempts to inform, clarify, or sell in normal language.

When doing this, keep in mind that Google understands synonyms, meaning that you don't need to repeat the same keyword too often. Google states that "cat breeder" will match for "kitten breeder"; "homicide" for "murder"; and "song words" for "lyrics".[4] Thus, feel free to use synonyms of your keyword throughout your page.

Remember too that the Googlebot approximates the importance of a term with how early within a HTML document that term occurs. On top of this, the Googlebot might only scrape the first X bytes of your page—so if your header or paragraph tags containing keywords appear deep within your web page, they risk being underweighted or even missed.

Carefully Craft a Unique Meta Description

The meta-description tag *suggests* to Google what it ought to display in your Google Search preview snippet.

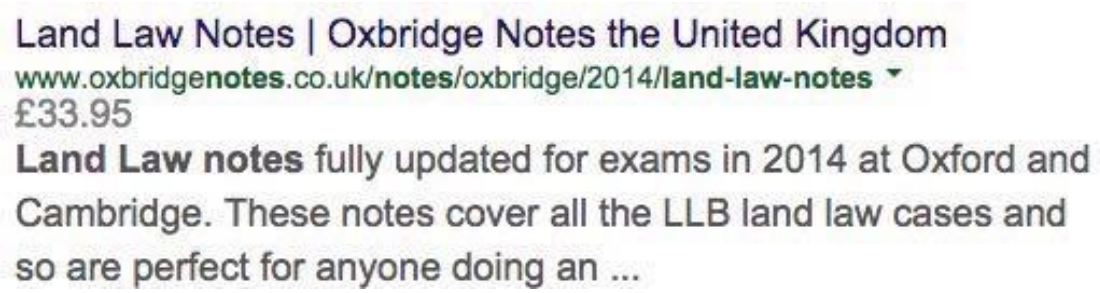


Figure 4.2 – *The Google Search preview snippet for one of Oxbridge Notes' landing pages*

The meta-description does not bind Google to the text you set; sometimes they will bypass your recommendation and create their own meta-description by picking out content from the body of your text that they believe to be a good fit with the search query.

[4]https://googleblog.blogspot.de/2010/01/helping-computers-understand-language.html

www.dbooks.org

I consider the meta-description tag as a sort of online advertisement. Google searchers see your search snippet alongside that of nine other competitors on the page, and it's the job of your meta-description to convince the searcher to visit your website instead of the others. That is the meta-description tag's sole purpose in life; it has no influence on how high you rank, only on how effective your snippet converts at its current rank.

Ensure Your Website Makes Sense to Googlebot

Google is best at indexing HTML and text content. Websites heavy on Javascript were, until recently, advised to recreate their Javascript content in HTML as part of a "graceful fallback", otherwise they would fail to rank. As of 2015, this is probably unnecessary, as shown by the experiments carried out over at Search Engine Land.[5] Just to be on the safe side though, Javascript-heavy websites should double-check that their pages are getting indexed by Google[6] and, on the off chance that their pages are being skipped, invest in graceful fallbacks.

Google's dealings with alternative content formats is hit-and-miss. PDFs are crawlable, except when they're not; for this reason, it's probably a good idea to convert them to HTML so as to enhance crawlability. Google cannot index video or audio content, so businesses that primarily produce these formats are advised to transcribe their work. Railscasts, a website that sells video tutorials about how to program, adds HTML transcriptions next to their videos. Mixergy, a website containing interviews with successful entrepreneurs, does the same, transforming their audio into text. (Sidenote: You need not actually transcribe yourself because there are services out there to do this for you.)

If you don't get around to converting content to HTML, the next best option is to optimise the text elements of these non-text formats. In the case of PDFs, this entails putting your keywords into

[5]http://searchengineland.com/tested-googlebot-crawls-javascript-heres-learned-220157
[6]Google Search Console > Crawl > Fetch as Google

the file's metadata, author field, description field, etc.[7]

There's a fairly common web design practice that's a big no-no for SEO: the display of some website text exclusively through images, without a textual backup. Although the text images might look prettier—indeed, that's usually why they are used in the first place over regular HTML—they have the downside of Google not being able to reliably extract and index the text within. Instead, you should display your text in HTML—and if you must use an image, be sure to attach an alt attribute or nearby transcription. (More on alt attributes in a later chapter.)

Include Rich Snippets and Structured Data

Google defines snippets as "the few lines of text that appear under every search result…designed to give users a sense for what's on the page and why it's relevant to their query." Rich snippets are simply search results that have been enhanced and enriched with nifty little extra titbits.

Google explains, "If Google understands the content on your pages, we can create rich snippets—detailed information intended to help users with specific queries. For example, the snippet for a restaurant might show the average review and price range; the snippet for a recipe page might show the total preparation time, a photo, and the recipe's review rating; and the snippet for a music album could list songs along with a link to play each song. These rich snippets help users recognize when your site is relevant to their search, and may result in more clicks to your pages."

At the very least, rich snippets make your search result entry look larger. All things being equal, the larger an entry, the more attention it commands—and the more attention something commands, the more likely it is to trigger clicks and visits.

Sold? Great! You can avail yourself of rich snippets by adding ad-

[7]https://moz.com/ugc/how-to-optimize-pdf-documents-for-search

www.dbooks.org

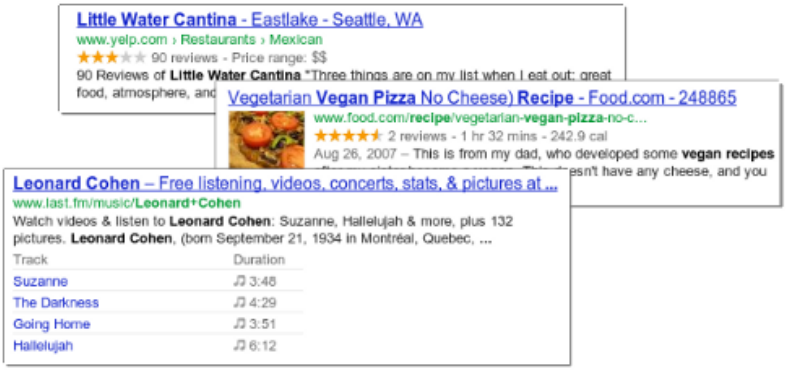


Figure 4.3 – *A selection of the special Google Search results facilitated by Rich Snippets. Source: Google.com*

ditional markup to your HTML that helps Google understand your content. Various formats are available for marking up your data, but Google recommends one called microdata, so you might as well follow their lead.[8] At present, Google supports rich snippets for content types such as reviews, people, products, businesses, recipes, events, and music. For a full and up-to-date list of supported formats, refer to the listing on Schema.org, a joint initiative by Google, Bing, and Yahoo! aimed at creating a set of agreed-upon schemas for structured markup on pages. By following these guidelines, your markup will be supported in other search engines too—at least theoretically....

Another great bonus accruing to some lucky search result entries are what Google calls "sitelinks". Underneath the entry for your search result, Google sometimes adds a collection of links to other pages of your website that they believe relevant.



Figure 4.4 – *The most basic search engine result without any sitelinks. Compared to what will follow, this is a rather boring showing!*

So where do you sign up for *your* sitelinks? Unfortunately, you

[8]https://support.google.com/webmasters/answer/176035

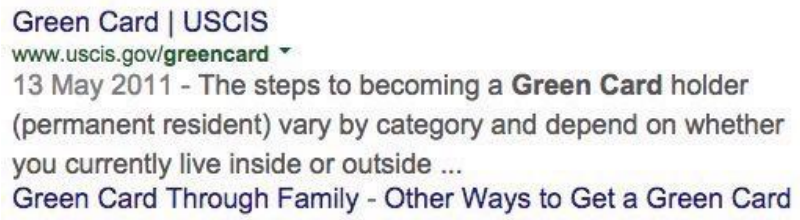


Figure 4.5 – *Notice the miniature sitelinks at the bottom. They don't consume much search result real estate, but they do serve to surface interesting content on your website and increase your search entry's click-through rates.*

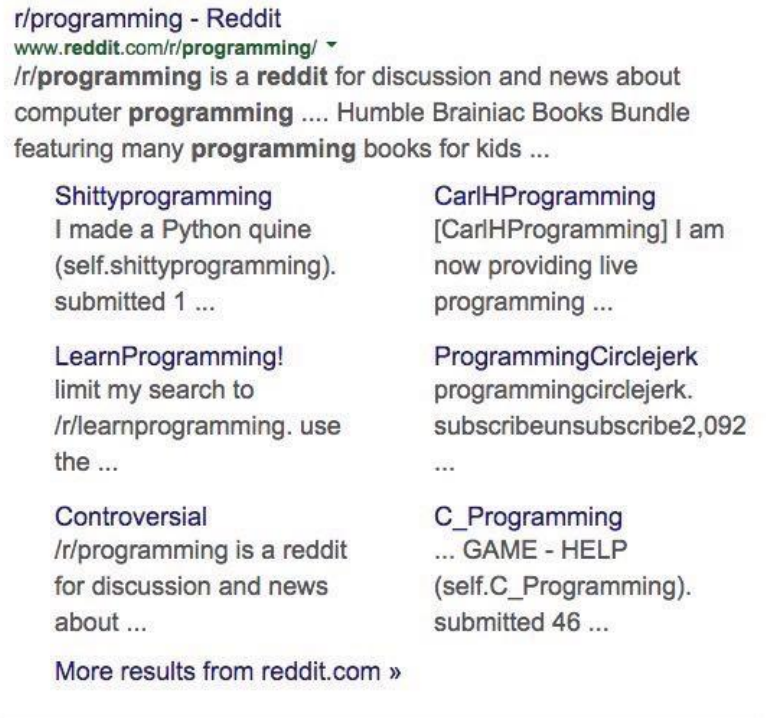


Figure 4.6 – *The glorious sprawl of a fully fledged sitelinks showing. I saw this when I searched for "Reddit programming". Before seeing this, I had never heard about /r/Shittyprogramming, but now I am intrigued—a win for Reddit's sitelinks!*

www.dbooks.org

cannot *explicitly* direct Google Search to promote chosen pages as sitelinks. Instead, Google's algorithm determines your sitelinks basis on each page's popularity with visitors, and with the ease they have in inferring your website's structure. You can help Googlebot out in this last regard by using the breadcrumbs microdata schema.[9] There's another lever of influence in here too: If Google happens to surface an inappropriate, sensitive, or incorrect sitelink, there's an option to demote individual links within Google Search Console.[10]

(Incidentally, you explicitly set sitelinks accompanying *advertisements* created with Google Adwords.)

Every website has one specialised content type that they focus on, be that tech concerts, dentists, or digital cameras. If you don't find your particular content type currently listed on Schema.org, Google recommends that you either:

- Use the Schema.org extension system to define a new type (and pray to God it gets ratified).[11]
- Use a less specific markup type. For example, Schema.org has no "Professor" type. However, if you have a directory of professors in a university department, you could use the Person type to mark up the information instead.

After adding the markup, you may be unsure about whether your efforts resulted in validly formatted data. To assay any doubts, refer to Google's Structured Data Testing Tool.[12]

Every Page Should Include a Picture

Researchers at OkDork.com analysed over 100 million articles and found that those which contained at least one image were twice as likely to be shared on Facebook as those with no im-

[9]https://support.google.com/webmasters/answer/185417?hl=en&ref_topic=4599102
[10]https://support.google.com/webmasters/answer/47334?hl=en
[11]http://schema.org/docs/extension.html
[12]http://www.google.com/webmasters/tools/richsnippets

ages.[13] This shows how important it is to augment your web pages with visuals.

Write Longer Pages

There are findings both by evidence-based marketers in the SEO world[14] and academic researchers[15] that web pages with higher word counts are more likely to be shared. And because sharing generates links and links generate further links, this is research well worth heeding. With that in mind, aim for at least 2,500 words per page.

[13] http://okdork.com/2014/04/21/why-content-goes-viral-what-analyzing-100-millions-articles-taught-us/

[14] http://okdork.com/2014/04/21/why-content-goes-viral-what-analyzing-100-millions-articles-taught-us/

[15] http://www.nytimes.com/2010/02/09/science/09tier.html?_r=1

www.dbooks.org

5 | SEO at the Whole-Website Level

Generate a Sitemap

When Google is unaware of the existence of some page on your website, then they cannot index it, and consequently you stand absolutely no chance of ranking for the content within. In short, its SEO potential is completely squandered. Not all pages are equally vulnerable to being overlooked by Googlebot. Pages linked to from other pages (on the same domain or on external domains) are fairly likely to be indexed. But pages only accessible through form submissions, on-site searching, or Javascript/Flash manoeuvring, are likely to be missed.

These problems can be overcome with a sitemap, a humble yet often overlooked tool for improving a website's exposure in search engines. It works by simply listing all the URLs that ought to be indexed, thereby alerting search engines to pages that might otherwise be missed.

You can submit sitemaps for alternative types of content too, such as images, videos, and even microdata (PageMap sitemaps).

Link Internally

Internal links are the selfie sticks of the linking world. They are the links you apportion to a web page from other parts of your website (i.e., links originating from the same domain). The bog standard example goes as follows: Your homepage (selling electronic devices) links to 10 product-category pages (e.g., "/cameras" and "/speakers"), and each of these product-category pages links to 20 product pages (e.g., "/speakers/ue-boom", "/speakers/bose-soundlink-mini"), which themselves complete the circle by linking back to both their parent product-category page ("/speakers") and to the homepage ("/"). This web of inter-

linking pages, intentionally configured by you, constitutes your internal linking strategy.

Internal linking strategies are useful for SEO because Google's algorithm perceives the web in terms of individual *pages*, and links from page to page on the same domain count as valid ranking signals. Sure, an internal link might not be as strong a signal as one coming from an external, third-party website, but it still counts for something. Therefore, you'll want to ensure that every page on your website is linked to at least once by another page.

This isn't a licence to overdo it with internal linking. There is a limit to how many links Google will crawl on a given page, roughly about 150. Moreover, Google views pages containing hundreds of links as sorts of seedy internet slums deserving of penalisation in their rankings. You can avoid coming across as overzealous in internal linking by sticking with a simple strategy, like the tree structure we saw above ("category > sub-category > sub-category").

The mere presence of internal links is only half of what is needed in your internal linking strategy; the other half involves getting the anchor text right. This is because anchor text counts as a ranking signal, so you'll need to write something meaningful, ideally containing your keyword.

Imagine I was linking to this very marketing book from another web page. Here are a few anchor text options I could have chosen, some better and some worse:

(The words in italics are the links.)

1. "learn more about my marketing strategy *here*"

   By anchoring only on the keyword "here", I throw my SEO juice down the toilet. No one in their right mind will be searching for the keyword "here", not least an audience interested in reading my marketing book.

www.dbooks.org

2. "read about *marketing-driven development*"

   This approach is better, in that I use descriptive anchor text. However, it falters in that "marketing-driven development" is a term I just came up with, and a typical member of my target audience couldn't possibly be searching something I just invented. This anchor text thus fails as a signpost to those who haven't yet encountered my ideas. (Of course, there is potentially a longer-term play if my strategy were to colonise and brand the term).

3. "read my *guide to online marketing*"

   At last, a sensible and conservative choice of anchor text. This works because I described what I was linking to in the language of everyday use, helping my link show up in random searches made by my potential readership.

Use HTTPS Encryption

When someone asks you to write out the full internet address into your browser's URL field, you begin with "http://". These letters mean that the website is reachable over a particular technological protocol known as Hypertext Transfer.

This technology, however, is insecure because it isn't encrypted. As a result, malicious hackers can easily read the data you send over HTTP. For this reason, more and more professional websites prefer to communicate over an alternative protocol, HTTPS, which has heightened security. HTTPS websites are accessed by typing "https://" into the browser, where the added "s" stands for "secure".

Google wants the whole web to use HTTPS, and, to this end, they officially announced that websites using HTTPS will receive a ranking bonus in their search results.[1]

[1]https://webmasters.googleblog.com/2014/08/https-as-ranking-signal.html

This means that you should ensure that you have HTTPS installed, and also that it's turned on for *every page* of your website (not just the credit card and login pages, as is default with many HTTPS installations). Because your old visitors may have bookmarked your website as "http://", you may also need to redirect any "http://" traffic to its "https://" counterpart.

www.dbooks.org

6 | SEO by Mitigating Duplicate Content Issues

Many websites have the exact same (or very similar) content viewable through the different URLs of their various websites. This is known as "duplicate content" in the SEO world, and it's something you want to avoid.

What's the problem with duplicate content? To start with, a website containing the exact same text repeated over and over on many different pages comes across as spammy, and search engines could penalise it for this reason.

Moreover, duplicate content causes all sorts of technical problems for search engines.

Firstly, search engine crawlers, even at the best of times, usually don't index *every* page on a website; instead, they make informed decisions about what content is worth keeping around and what content should be chucked aside. When the crawler encounters duplicate content, it (sometimes) takes this as a signal to ignore all but one of the duplicates. But how does the crawler pick which version to keep? That decision isn't easy for a machine, and the website owner risks the algorithm indexing some odd-looking, confusing version of the web page instead of the perfectly gorgeous page intended for human readers. Instead of putting their best foot forwards, the webmaster with duplicate content is instead presenting a smelly sock riddled with holes.

Secondly, if, as may happen, more than one version of the duplicate content happens to get indexed, search engines don't know how to apportion ranking authority across these duplicates, leaving webmasters open to the undesirable possibility that said ranking authority is split across all the duplicates and diluted so far that the webmaster can no longer compete with rival businesses

who have focused their ranking authority into a few no-duplicate URLs.

Duplicate content issues crop up if your website has two or more distinct URLs that point to the same page (or wad of content). This sounds simple, but the subtleties of this idea can get lost on casual web users. Specifically, you should worry about duplicate content if any of the following apply to you:

- Your website exists in cloned form across multiple domains and subdomains. The most common case of this happening is when your website responds to and serves up the same content for both "www.mywebsite.com" and "mywebsite.com". This essentially creates duplicates for every page on the entire domain—the only difference being that one version has "www" in front and the other doesn't.

- Your website has multiple URLs for accessing the same content. For instance, "/t-shirts" and "/category/t-shirts" might point to the same page (i.e., to the one that displays the t-shirts you offer for sale). This problem often occurs on the home page, which, on old-fashioned websites, can be accessed variously with "/index.php", "/home", and the root URL, "/".
- You append parameters to URLs, and some permutations of these parameters lead to the same or very similar content as other permutations. For instance, sort parameters, such as "/t-shirts?sort=price_asc", lead to the exact same content, only rearranged. Sometimes filters parameters lead to duplicates too, such as when "/t-shirts?review_score_greater_than=3" (which filters a page to only show items with review scores over 3) acts on a data set in which every item has a review score of least 4 anyway, meaning that the filtered page is exactly the same as the unfiltered original. Search engines don't automatically strip these parameters out, as you might expect.
- Your website responds with a HTTP status 200 to both "/products" and the equivalent with a trailing slash, "/prod-

www.dbooks.org

ucts/". Surprisingly, Google treats both as distinct URLs, explicitly stating as much in their official webmaster literature.[1] They mention that this practice is "often OK" but not "perfectly optimal". A careful webmaster should ensure that their website only responds to one of these slash possibilities.
- Your website shows the same content to both uppercase or lowercase versions of a URL (e.g., "/taxons/t-shirts" and "/taxons/T-SHIRTS" both lead to the same content).

Just to complicate issues even further, there are cases when you *specifically want* the duplicate (or near-duplicate pages) to appear within Google Search results. Most often, this is with translated/regionalised pages that you'd like to put forward as local organic entry points to your website, to be shown or hidden in the results depending on where the searcher is based and in what language they are searching. In this specialised case, please ignore the ensuring tips for avoiding duplication and instead look to the hreflang directive, as elaborated upon in the chapter on SEO internationalisation.

Now, back to banishing *undesired* duplicate content.

Redirect Instead of Responding

This should be your first line of defence against duplicate content, owing to its ease of use and its low impact on SEO link juice when compared to the nuclear option, the robots.txt directive (which we'll see later).

Let's say you want your customers to be able to visit both "www.mywebsite.com" and "mywebsite.com". Instead of displaying content for (and responding 200 to) both possibilities, you should tell your server's router to respond with a HTTP 301 response code (permanent redirect) on one of the two possibilities, redirecting all those requests to your now officially sanctioned choice. The web page to which the redirect points will receive at

[1] http://googlewebmastercentral.blogspot.com.br/2010/04/to-slash-or-not-to-slash.html

least 90% of the redirected link's ranking power.

While we're on the topic of redirects, let's highlight another of their purposes in SEO: preserving link juice on retired or edited URLs. As you know, any given page on your website can potentially accumulate reputation, links, and Google ranking in organic searches. But whenever you edit its URL name (or drop the page altogether), you reset the counters at zero, throwing away every last drop of the SEO juice earned by that URL. In these circumstances, the 301 response code comes to the rescue. Its effect is a memo to Googlebot saying, "Hey Googlebot, the old page you're looking for has relocated. Please update your index and transfer all my SEO ranking accordingly. Kthxbai." This point is important. Indeed, I would go so far as to say that editing a URL name without placing a 301 permanent redirect on the old one is as about close to SEO suicide as I can imagine. Yet it happens all the time.

It is tricky to ensure that redirects are always created after editing old URLs. Not only do you need to watch out for programmers who edit the URL structures, but also admin staff who edit links in the CMS and regular users who change their own content, as they so often tend to do. To meet these challenges at a systematic scale, consider implementing non-changeable permalinks attached to each piece of database content, thereby freezing the original URL and immunizing it against change. If this isn't possible (say because it is necessary for your customers to change URLs), consider a database-backed system that automatically remembers your old URLs and responds to requests for them by 301-ing them forward to their latest incarnation.

Use Canonicalization

Unlike with redirects, with canonicalization your server will continue to respond (with HTTP status code 200) to the duplicate URLs. Instead, there will now be snippets of HTML added to the duplicate pages, and the purpose of these is to explain to Google which page is the "canonical" one (i.e., the definitive one that

www.dbooks.org

should show up in their search results and accrue all the search juice).

Here's an example of canonicalization in use, as you might see on a page "http://www.mywebsite.com/index.php" which duplicates the root URL.

```
<link href="http://www.mywebsite.com" rel="canonical">
```

Canonicalization is a useful advantage when you want to have distinct versions of your content for other people to bookmark or link to (e.g., you have distinct pages for printer-optimised alternatives for your content).

Just as with redirects, canonicalization preserves the link juice of the duplicate pages.

The href attribute of the canonical URL tag accepts fully qualified absolute paths (includes "http://") and relative paths (doesn't). I recommend sticking to fully qualified absolute paths, since poorly formed paths can potentially lead to insidious, damaging errors (see Google's warning on this topic[2]). For example, if you write "mywebsite.com/t-shirts", it will attribute your SEO juice to "http://www.mywebsite.com/mywebsite.com/t-shirts"—a non-existent page. This happens because the initial "/" was missing in the relative path.

By the way, there's a potential gotcha when using absolute paths: You need the proper protocol ("https://..." vs "http://...")

Write a Robots.txt or Meta Robots Tag

The robots.txt file, which should be retrievable at the URL "/robots.txt" of your website, is a directive asking web robots—in particular search engine crawlers—not to crawl specified pages on your website. This can be thus used to stop crawlers indexing duplicate URLs.

[2]http://googlewebmastercentral.blogspot.com.br/2013/04/5-common-mistakes-with-relcanonical.html

Despite this ability, I would nevertheless recommend against employing the robots.txt for these purposes, and I instead advise you to reach for redirects or canonical URLs (or the meta robots tag...).

Why so? Duplicate URLs sometimes accrue decent SEO juice (e.g., the mobile-friendly version of your website ["m.mywebsite.com"]). Ideally, you want to channel this SEO juice into your main page (e.g., "www.mywebsite.com"), so as not to lose any authority. The problem with listing URLs in the robots.txt file is that you lose ALL their SEO juice. As Adam Audette puts it, it is a "Pagerank dead end", "a sledgehammer".[3] This waste doesn't happen with other duplicate-content mitigation mechanisms such as redirects and canonicalization, which are both much more effective at preserving SEO juice.

All that said, there *are* good use-cases for the robots directives. Principal amongst these is removing low quality *non-duplicate* content from Google's index. At first it might seem counterintuitive to want to *reduce* your website's crawlable surface area, but for many webmasters this approach pays dividends. Having a robots.txt restriction stops crappy entry points to your website from appearing in your search engine results. For example, at my company, Oxbridge Notes, we have an (admittedly crappy) URL design whereby every product page links to a nested page for sending questions to the seller of that particular product. That gives us URLs like "/land-law-notes/buyer_inquiries", "contract-law-notes/buyer_inquiries", and so on, thousands of times. I didn't want these content-dry generic question-form pages littering up my search results and competing with my optimised landing pages (the parent product pages, such as "/land-law-notes"), so, to this end, I updated my robots.txt notice with a directive not to crawl these "buyer_inquiries" pages.

It's true that I probably shouldn't have designed the website with this crappy structure, but, at this point, there's no point in fixing something that isn't too badly broken, especially when a quick entry to the robots.txt file dissolves away the issue.

[3]http://www.rimmkaufman.com/blog/robots-txt-best-practices-for-seo/03072012/

www.dbooks.org

Another good use-case of robots.txt is privacy. In the country where I reside, every website owner is required to place their contact details on their website on a page known as an impressum. As an individual running a web business with remote workers, I don't have a dedicated office. In this case, the law requires me to place my home address on the impressum page. This makes me uncomfortable about both my privacy and my safety. Not too long ago, a nasty mugger attacked me and threatened that if I report him to the police (which I subsequently did), he would wait for me outside my apartment (which, mercifully, he subsequently did not). This threat had teeth though, because I knew that all it would take for such a malicious character to find out where I lived was to google "jack kinsella address". (He knew my name from the ID card he stole from me.) The robots.txt could have, perhaps, reduced this risk, say if I instructed search engines not to index the page containing my address.

Robots.txt can also preserve privacy in less dramatic circumstances. For example, some webmasters leverage it to hide content and features not yet officially launched.

If you really want to restrict crawler access to content with the robots directives, I would advise you *not* to rely on the classic robots.txt file, but instead to rely upon the newer—and in my opinion superior—meta robots tag.

The big problem with the robots.txt file is that it doesn't quite do what it suggests it does; it only instructs search engines not to *visit* the URL. Surprisingly, search engines will nevertheless index the *existence* of that URL and display ugly, stripped-down, contentless ghost entries for search results. You probably don't want that:

These shallow results present a poor image and poor user experience to potential customers on Google, and, moreover, they might leak information to hackers or competitors—information that you'd otherwise wish to keep private, such as URL entry points or revealing parameters nested within URLs.

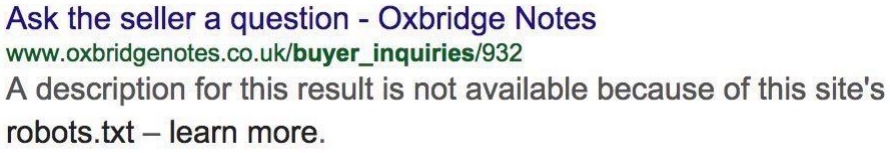


Figure 6.1 – *My robots.txt file contained the entry "/buyer_inquiries" and, contrary to my expectations, Google indexed thousands of these worthless stubs.*

In order to circumvent these limitations, a webmaster can insert meta robots tags within the HTML head of relevant pages, eschewing the need for entries in the robot.txt file. These tags accept multiple arguments, which allow for more flexibility in directing Googlebot how to treat each entry.

```
<meta name="robots" rel="noindex, follow">
```

The above example shows a meta robots tag instructing Googlebot not to index a page, but to "follow" (i.e., attribute link juice to) any links on that blocked page, meaning that external links to this "roboted" page won't be completely wasted as they would be if we had used a robots.txt entry. (Thus, this "follow" directive essentially undoes the major link juice disadvantage of regular robots.txt entries.)

For a full description of the various arguments accepted by meta robots tags, consult Google's guide on the topic.[4] (Interesting example: The instruction "noarchive" directs Googlebot not to include a page within Google's cache, a nice boon for privacy).

Last point: The meta robots tag is a property of HTML. What if you want to restrict the indexing of non-HTML content such as PDF files? The HTTP X-Robots-Tag header fulfils this role and accepts the same arguments as the meta robots tag, affording you the exact same feature set for content in alternative formats.[5]

[4] http://https//developers.google.com/webmasters/control-crawl-index/docs/robots_meta_tag
[5] http://googleblog.blogspot.com.br/2007/07/robots-exclusion-protocol-now-with-even.html

www.dbooks.org

7 | SEO in Image Search

It's increasingly important to conceptualise SEO not as one monolithic search results race for text, but rather as a series of events for various formats. Search engine optimised textual content enters you into the main event, but a webmaster might also choose to compete in Google's alternative races, such as Google Images, Google Videos, or Google Maps.

Given that these sidetracks may be undersubscribed by your competitors, you potentially stand to gain an advantage through your focus on the unconventional. The question you should be asking is what steps you can implement as a to boost your website's rank in these alternative searches.

Considering that nearly all modern websites use images, image search is probably the most common side event, so let's use Google Images as our example. While it's true that Google increasingly uses AI algorithms to recognise objects in images and match searches in this way, you can still exert influence by taking the following steps:

Output Public-Facing Images Through Your Software

Do your users upload images? If appropriate, make these public. What if they don't? Consider generating images out of other content. In the screengrab below, you'll notice that the top five results for one of my valuable keywords, "land law notes", are all image representations of the documents uploaded by my authors. Lacking any real product pictures for these electronic documents, I instead have an algorithm that photographs the first page of each document. This allows the customer to appreciate the formatting of the document without having to download the PDF, and it has the nice added benefit of ranking well in Google Images.

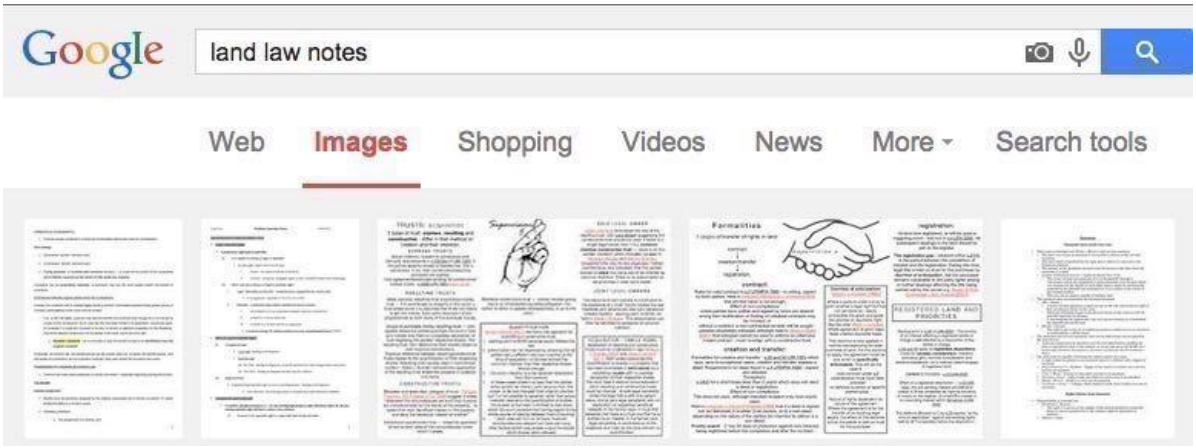


Figure 7.1 – *Here, you see enticing pictures of the documents on sale*

Add Alt Text to Your Images

The alt text directive has two big purposes. First, it provides accessibility for readers with vision difficulties. Second, it shows a text-based description of the image when a connection isn't fast enough to fetch an image. For the time being—a time before Google's object recognition algorithm has flourished—setting an alt tag on images works surprisingly well as a ranking signal, so don't forget to add it.

Choose Informative Filenames for Your Images

Many websites have auto-generated filenames for images (and other assets) that are incomprehensible to humans. If you instead write a descriptive filename containing your keyword, then you stand to rank higher in search results and to appear whenever users search specifically for files of certain formats in Google's primary search. It's not a stretch of the imagination to assume that "british_shorthair_cat.png" will rank higher than "01231231312.png" in a search for "British shorthair cat pics".

Write Captions to Accompany Your Image

In their image publishing guidelines, Google specifically mentions that they scan captions and other accompanying text for

www.dbooks.org

clues about how to rank your images.[1]

Generate and Submit an Image Sitemap

Image sitemaps, like normal sitemaps and video sitemaps, ensure that Google finds content on your website that they might otherwise miss. It's trivial to generate these sitemaps with open-source plugins, so you've no excuse not to.

[1] https://support.google.com/webmasters/answer/114016?hl=en

8 | SEO Through Backlinks

Backlinks are arguably the most important aspect of SEO, and indisputably the most frustrating aspect to implement, owing to the challenges and uncertainties inherent in convincing others to link to you and the shamelessness needed to pull it off. Even world-class SEO pros expect to have to solicit ten website owners for every backlink they secure; if you're a beginner, be prepared to experience a lot of rejection.

There are two ways in which backlinks win you traffic. The first of these is by raising your rankings in Google Search for your keywords, such that a higher proportion of searchers end up learning about and visiting your website. The second way backlinks win you traffic is by channelling a flow of visitors from blogs, forums, and other websites where people in your market assemble, such that these visitors travel your way even without having interacted with a search engine.

My reason for distinguishing these two modes of backlink traffic is because the second form is hugely underrated. Websites acting in some internet sectors, such as pornography, receive hordes of traffic despite being completely banned from appearing in Google Search.

In the eyes of search engines, there are normal links—which attribute search ranking juice—and no-follow links—which don't. In the olden days of the Web 2.0, it was possible for someone to link to themselves within the comment section on some other website's blog and gain a search engine boost because of this. But this practice quickly led to a massive rise in comment spam, and, as a reaction, webmasters started installing software that automatically transformed all links in blog comments into their no-follow equivalents.

www.dbooks.org

Many professionals working in search engine optimisation do not bother collecting no-follow links. But this restraint is misplaced if the no-follow link could have appeared on a page visited by millions of readers every month. While it's true that these links won't win any traffic via Google Search, they can (potentially) deliver thousands of visitors via referrals from that blog. Don't underestimate this.

There's another preliminary point worth making: Not all pages on your website are equally likely to attract backlinks. For example, if you sell life insurance, it's unlikely that any journalist or blogger in their right mind will ever link to the page detailing your pricing. But the same people wouldn't think twice about linking to an interesting article you post about the wacky conversations of a quirky door-to-door insurance salesman. It's unrealistic to SEO *every page* on your domain, so focus on the ones that are easy sells and let their inbound links boost the average rankings of the other, more boring pages on your domain.

Finally, you should budget, as a minimum, at least as much time and money into gathering backlinks as you did on writing the page you hope to rank for. That means that an article that took 20 hours to write deserves at least 20 hours of promotional efforts. If you aren't prepared to do this, then you might as well forget about having an SEO strategy.

Include Contributions from People with Big Audiences

This strategy is timeless and will be relevant for as long as there are people who wish to promote themselves. The basic idea is that you either a) solicit statements, opinions, or articles from others for inclusion in your website; or b) write about others in glowing terms.

Once you publish the page concerning them, these third parties will be motivated to promote it because they know that your success feeds back into their own. On a (now defunct) blog, I asked 20 bloggers in my niche to give their top advice for the new year. When this post was published, it won me more traffic in one day

than all my other posts put together had earned in the past year.

And, by the way, if you follow plan b) (writing about others in glowing terms), remember that these people will almost definitely not be aware of your work unless you tell them; you'll need to reach out and let them know when you publish.

Coin Concepts

Have you heard of the email management technique Inbox Zero? Or the term for a passive income business, Muse? Or the Suzuki Method for learning music? These are all relatively recently invented concepts that stuck and have entered our vocabulary.

The advantage of coining a concept is that people will always have to refer back to you (via a backlink) when they write about it or explain it elsewhere. Done tastefully, this is a great tactic.

Leave Your Link Where You Can

For a start, you should share your link on your (public) social media profiles and status feeds. These links no longer contribute towards Google rankings[1], but are still worth having so as to reach non-search-engine referral traffic and possibly convince them to take a gander over at your website.

Next up are links left on social bookmarking websites like Reddit, Delicious, or StumbleUpon. If your writing is up to par and your topic sufficiently interesting, it's not at all unrealistic to gain decent traffic (and Google rankings) through a post to the appropriate topic-based section of these websites. If possible, have a friend submit on your behalf—the moderators of many communities dislike self-promotion.

The next stop on your self-posting crusade is commenting—be that on blog articles, forum posts, social bookmarking websites, or question and answer websites (like Quora or Stack Overflow).

[1]https://searchenginewatch.com/sew/news/2325343/matt-cutts-facebook-twitter-social-signals-not-part-of-google-search-ranking-algorithms

www.dbooks.org

A phenomenal source of traffic, if you can find it, is post-it-yourself communities of like-minded people. For example, in the programming world there are a bunch of community-maintained news websites that collect links to the hottest new articles in that field. These can be a great way to reach a large, targeted audience without having to get past the hurdle of an unpredictable content curator.

Email Outreach

This (rather cheeky) approach consists of cold emailing people in your niche and telling them about your article with an email that says something like, "Love your blog, esp. the article on X. I thought you'd like this article I wrote on the same topic."

Don't ask them to share or give you a link in your email—but don't be surprised if they do.

Last time I tried this approach, I got about a 33% promotion rate (which I define as receiving at least a Tweet).

Pitch Your Content To Writers of Round-Up Articles

A perennial favourite among internet users are articles with titles like "Complete List of Twitter Clients for Android" or "99+ Best Web Design Resources" or "Portable Bluetooth Speaker Round-Up 2016". Considering that the purpose of these articles is to enumerate what's available, it's not difficult for you to convince their authors to include whatever you're offering.

Prominent members of the SEO community, such as Brian Dean, agree that your odds of being included in the list are improved if, in your first contact with the author, you point out something outdated in their article[2] (e.g., one of their link 404s, a company they mentioned has rebranded, or the prices they list for a product are obsolete). You have done them a favour, and they owe you one back—payment for which you accept in the form of backlink.

[2]Brian Dean goes into this in depth and suggests techniques for doing this: http://backlinko.com/high-quality-backlinks.

Where this tactic isn't possible, it's still worth reaching out to these authors anyway. When doing so, it is recommended that you avoid sending your link to the author in your first email; instead, you should feel out their interest and ask them if they are OK with you forwarding an article for their opinion. Only if the author shows some interest should you send the goods.

Piggyback off High-Ranking Websites

It's very possible that a page on the third-party website which links to your website will nevertheless outrank your website. Here's a shady example of this idea in action: I once googled "best indian in Berlin" and noticed that the top result was a link to a Yelp page with a review containing the text "best indian in Berlin". I subsequently ate there, and it was dreadful. When I returned to Yelp to poke around, I noticed that the reviews were obviously faked (e.g., accounts with no other activity, all created on the same day). The restaurant simply figured out what potential customers typed into Google, found a way to insert that text into a website with decent Google rank, and then benefitted accordingly. Had they written the same text on their own website, it would never have ranked and come to my attention. But by placing a link to themselves on Yelp, this horrid little restaurant was able to leverage Yelp's popularity.

Guest Post on a Large Platform

The idea is that you write something of extraordinary quality and give it away for free to someone who owns a platform (e.g., to a large blog or a major news website). The platform wins by getting good content for free and you win by slipping a backlink into the copy you send them.[3]

[3]Specific strategies for finding guest blogging opportunities and pitching are detailed here: http://backlinko.com/the-definitive-guide-to-guest-blogging.

www.dbooks.org

(Maybe) Buy Links

Google warns against doing this in the strongest of terms, but their bark is a substitute for their inability to enforce this policy, at least when links are bought covertly.

Let's start with what NOT to do: Don't buy links in a link-buying marketplace. If it's public knowledge that a domain sells links, then Google could theoretically detect the sale and penalise you for breaking their policy.

If you want to get away with buying links, the key is to contact a webmaster established in your niche and do a private deal. The link they give you should be integrated naturally into the third-party website's content such that your paid-for link is indistinguishable from any other natural link. With these precautions in place, there is little chance of Google detecting your bribery.

9 | SEO Vis-à-Vis Social Media

Web pages with higher shares on Twitter/Facebook or more upvotes on Reddit enjoy higher search rankings. It is unknown whether this is because Google directly reads the share/upvote counts and attributes search juice accordingly or because shares have secondary benefits, such as referring more traffic, some of which will eventually backlink to you from elsewhere (e.g., by mentioning you in their blog).

Either way, it's important to encourage people to share your stuff on social media. For a start, you'll want some tasteful "share" buttons. But these alone aren't enough; you should also optimise how your website is presented when it appears on social media. What image accompanies the link to your website? With what description and title text does it appear? By controlling these factors, you can optimise second-level shares (i.e., how often an existing share/tweet gets shared again).

Include Open Graph Tags

If you've ever been frustrated about Facebook's poor preview photo or website title choice that shows whenever someone shares one of your pages on Facebook, then you'll like Open Graph tags.

Here's how they work: Whenever someone shares your page on Facebook, the Facebook software sends out a crawler whose job it is to generate a preview of your web page for other Facebook users. If available on your page, the crawler parses something called Open Graph tags, a special format which has parallels with microdata, Twitter Cards, and even regular old-fashioned SEO classics like meta-descriptions.

Open Graph is largely associated with Facebook, but other so-

www.dbooks.org

cial media websites—such as LinkedIn—also scoop up this data, rewarding you with some extra bang for your buck when you implement these tags.

```
<meta property="og:title" content="The Rock" />
<meta property="og:type" content="video.movie" />
<meta property="og:url"
      content="http://www.imdb.com/title/tt0117500/" />
<meta property="og:image"
      content="http://ia.media-imdb.com/images/rock.jpg" />
<meta property="og:description"
      content="Directed by Michael Bay. With Sean Connery,
      Nicolas Cage, Ed Harris, John Spencer." />
```

The above depicts an example of some Open Graph data from everyone's favourite movie database, the IMDb. Open Graph data gives you a degree of freedom independent of the HTML, facilitating the variance of your Facebook presentation independently of your Google presentation, and letting you modify your message, if needs be, for your sundry audiences.

Be aware that Facebook has a "no-pixelation policy". They expect a surprisingly large image (1200 x 630 pixels at time of writing), at least if you want your Facebook share post to be accompanied by a banner-sized image. If, instead, your Open Graph data only links to a low-res pic, Facebook won't scale it up, meaning that you'll be left with a dramatically less impressive and radically poorer-converting miniaturised picture.

Check out the difference between these two image types with this pair of pics taken from Facebook's official documents.

Last note: Facebook offer a tool for debugging your Open Graph metadata so as to ensure it validates.[1]

[1]https://developers.facebook.com/tools/debug/

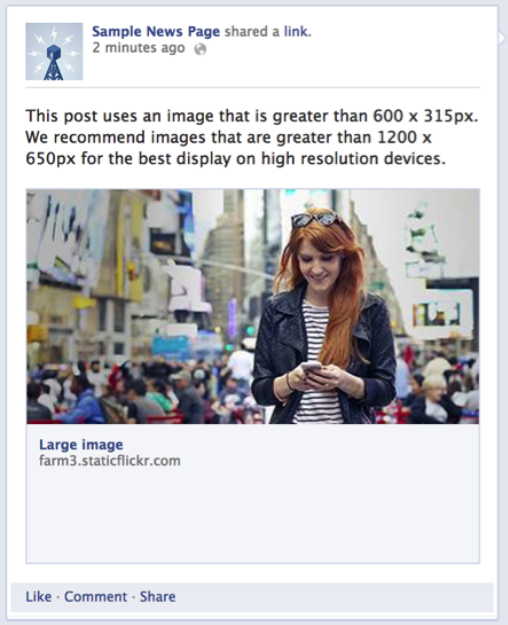


Figure 9.1 – *The high-res banner pic*



Figure 9.2 – *The low-res pic*

www.dbooks.org

Include Twitter Cards

Your customers and fans will, I hope, tweet about your product. If you wish to maximise the benefit you'll receive from these spontaneous tweets, look at Twitter Cards. These are special sorts of tweets that add extra functionality whenever someone clicks or expands the basic tweet. Twitter Cards can add playable videos or songs, article summaries, photos (or photo galleries), download links to apps, specifically designed call-to-action buttons (called the "website card"), or even miniature sign-up forms. To avail yourself of these possibilities, you need only add the appropriate metadata to the head section of your web pages. Twitter will parse this data and create a Twitter Card for you on this basis.

If you're uncertain about the correctness of your website's implementation of Twitter Cards, test it with Twitter's Card Validator.[2]

Promote Your Content on Social Media More Than Once

Many website owners announce their content on social media once and once only. This is suboptimal. Some of your followers will be offline when you release, and, for this reason alone, they will fail to see it. You can overcome this problem (and squeeze a lot more from your content) by posting the link a couple of times, starting with a week later.

There's a more sophisticated approach to this tactic: If you conceive of your content in terms of a list of tips (or quotes or ideas), then you can tweet a different one of these tips every day for the next few weeks. Software services, such as Buffer, even enable you to schedule social media in the future.

[2]https://cards-dev.twitter.com/validator

10 | SEO and Internationalisation

So you've seized the market in your homeland and now it's time to build your empire abroad. Enter the need for internationalisation. This concept, alternatively known as regionalisation, encompasses all the changes you'll need to make in order to serve customers in your newest regions. This includes translating text into other languages, displaying prices in other currencies (taking into account foreign tax rules), switching time zones, selectively showing or hiding content depending on regional relevance, and installing new payment and delivery options.

Nearly all of these internationalisation modifications will change the public surface area of your website, which means that they will also impact your website's SEO. As we have seen in earlier chapters, Google dislikes duplicate or near-duplicate content. This can be a problem, because internationalisation efforts such as minor regional modifications run the risk of triggering these duplicate-content alarms and thereby thwarting SEO efforts.

This chapter helps you steer through these straits.

Distinguish Language from Territory in Architecture

Let's start by disentangling two concepts in internationalisation that tend to get lumped together: language and territory.

Think about this: It is possible for a website to be focused on the *geographic region* of, say, the United States in terms of product range, delivery possibilities, and time zones, but nonetheless want to serve this territory through various *languages* (e.g., in English and in Spanish, the two main languages spoken in the US).

In a similar vein, it's also not difficult to imagine a website whose English language pages could span many geographic regions,

www.dbooks.org

serving various English-speaking countries like Ireland and New Zealand—and also serving expat contingents resident in other countries (e.g., in Germany).

I bring up this distinction because a website architecture ought to be configurable to modify regional and linguistic content independently of one another, instead of merely conflating the two concepts together. If one were to categorically say that "this is our American site" or "this is our German site", this would be a failure to appreciate the more subtle demands of going global.

Choose a Fitting Structure for Internationalisation

After deciding that it's time to turn a previously one-country-only website into a more cosmopolitan affair, the biggest question from the SEO perspective is that of how to structure the URLs to accommodate all this newly internationalised content.

One temptation might be to retain exactly the same URL structure as you had pre-internationalisation, and instead use cookies or Javascript to display translated versions of the text. With this setup, the URL "example.com/about-us" would be served to both German and American visitors, yet readers in each of these regions would view differing content specialised for their particular regions.

Don't do this. Google strongly advises against it, specifying that regionalised content ought to appear on *separate* URLs. Search engines expect that each URL has an unchanging identity; they are unprepared to probe a page by tweaking language and region variables, so internationalisation efforts that deploy these technologies inadvertently have the effect of hiding the internationalised parts of the website from search engines, which is not what you want.

Their condemnation also makes sense when we think about the user experience perspective:

1. Why should we expect a German speaker to understand the

URL "/about-us" when the equivalent in their language is /ueber-uns? This sort of presumption would amount to linguistic colonialism.

2. The user experience of *automatic* internationalisation is maddeningly frustrating. What if an American visitor borrows his German friend's laptop and wishes to place an order in English rather than in German? Is he compelled to fiddle with browser language settings and rummage for some other kind of hack just to reach the English content? Surely it would be a better user experience if the internationalisation were obvious from the URL structure alone, enabling the visitor to switch to their preferred regionalisation without the muddle (e.g., by changing "de.example.com" to "en.example.com").

Now that we've decided that it's better to place internationalised content on separate URLs, which particular URL structure is best suited for these purposes? Let's go through the various options one by one and evaluate them:

1. URL parameters: "example.com/?country=fr", "example.com/?country=de". We're going to throw this one out immediately because Google specifically advises not to internationalise in this way.[1] Their primary reason is based on difficulties with URL segmentation (i.e., it is challenging for high-speed computer programs to reliably discern the different countries by scanning the URL parameters...the data just isn't as structured as it is with some of the other options below). And just as this structure is challenging for Google's high-speed indexers, so too will it be a nuisance for you (e.g., when working with segmentation in Google Analytics or other tools that work with URLs). I can also see a third flaw with this approach: *Country* is one of the most commonly used GET parameters for on-site searches or filters (e.g., a filter to show the book retailers whose country=Spain). Dedicating this commonly used parameter to internationalisation risks clobbering their other URLs.

[1] http://support.google.com/webmasters/answer/182192?hl=en&ref_topic=2370587

www.dbooks.org

'Nuff said—next please!

2. Country specific domains: "example.fr", "example.de", "example.co.uk". This approach is, without a question, crystal clear to visitors and search engines alike, and it also has the most professional air to it. That said, I would advise small website owners to steer clear of this approach based on the following arguments: Firstly, country-specific domains are heavy on administrative and other "busy work". Many countries, such as Australia, require you to present their authorities with a case for why you should have some particular domain, a process that could necessitate annoyances such as registering a trademark or business in the region. Secondly, the availability of your desired local domain names is not guaranteed, either because of the presence of other businesses or because of opportunistic domain squatters who snap up variants of your domain and hold them ransom to you. As soon as there is *just one* territory with a domain name unavailable to you, the technical and branding purity of this "one country-specific domain for each country" approach is irreversibly tarred because you will now need to resort to hacks and bandages in order to reach people in some countries. Thirdly, this approach is more expensive that the others, with costs accruing in terms of renewing domain names, SSL certs, and applying to governments for trademarks. Fourthly, this approach is inferior for SEO because it splits SEO juice across many domains instead of concentrating it into the one. Each of these domains will need to gather its own backlinks and reputation—hardly ideal.
3. Subdomains: "de.example.com", "fr.example.com". This approach is easy to administer and cheap to set up. Additionally, the website's internationalisation structure is crystal clear to the end-user. Google claims that subdomains are every bit as efficient as subdirectories for imparting SEO juice, but experiments by SEO experts show that there is still a significant penalty for subdomains over subdirectories (which we'll detail next).[2] The crux of the issue is that

[2]https://moz.com/blog/subdomains-vs-subfolders-rel-canonical-vs-301-how-to-structure-links-optimally-for-seo-whiteboard-friday

Google sometimes views separate subdomains as separate websites, meaning that they don't always pool the ranking power from all the subdomains together, which has the effect of dulling the SEO power. Subdomains are best avoided for this reason.

4. Subdirectories: "example.com/fr/", "example.com/de/". This approach is easy and cheap to administer, since there is only one domain to buy and deal with. It is also the easiest to program because a website partitioned by folders is about as standard as it gets, and doesn't necessitate complex fiddling with host files, SSL certs, or DNS configurations. Google's SEO meister Matt Cutts has praised this approach over subdomains because internationalising with directories leaves open the possibility of hosting separate services on your various subdomains ("maps.google.com", "mail.google.com") instead of filling up those slots with internationalisation information.[3] And, more importantly, it's seen to have the strongest SEO effect. In short, this is the approach for you.

Communicate Your Structure to Google

It is not enough to just internationalise your website; you also have to clarify to search engines what your internationalisation structure is. This serves two rather important purposes: Firstly, signalling your internationalisation structure helps steer you free of Google's penalties for duplicate content. Secondly, it helps ensure that regional searchers see appropriate content from your website (e.g., you don't want googlers based in Italy to see search results for your UK website—this mismatch would scare off organic traffic).

The Hreflang attribute

This is a special piece of data added to parts of your website in order to inform Google that regionalised alternative versions of content exist.

[3] https://www.mattcutts.com/blog/subdomains-and-subdirectories/

www.dbooks.org

There are three different ways of specifying this hreflang attribute: within HTTP headers, within the HTML head, and within your sitemaps. Having these three possibilities open to you doesn't, however, mean that you should use any more than one of these formats—besides this being redundant, this sort of repetition raises the potential for errors and inconsistencies.

The hreflang attribute does not need to be set for *every* page of your website. Instead, it is only necessary for pages that appear in two or more regionalised alternatives. Let me illustrate this idea by referring to an imaginary pharmacy website. This pharmacy sells the drug paracetamol in all three territories it serves (France, the UK, and Russia), and the language, currency, and delivery options vary depending on the territory in which a visitor resides. Because there are three regional variations of this paracetamol page, it makes abundant sense to attach a hreflang attribute to it.

But the pharmacy also stocks other drugs—for example, controversial narcotic painkillers that are illegal in France and the UK. These medicines only appear on the Russian website. Because there is no mirrored content in any other region, there is no need for a hreflang attribute to be specified in this case—indeed, doing so would send Googlebot down a dead end.

Let's dig a little into the specific format of the hreflang attribute. Here's what you would expect to see on a page that has three international variations:

```
<!-- This page is example.com/en/about-us, where
the programming team chose to put the href
attribute in the HTML head. -->

<head>
  <link rel="alternate" hreflang="de"
        href="http://example.com/de/ueber-uns" />
  <link rel="alternate" hreflang="en-GB"
        href="http://example.com/en/about-us" />
  <link rel="alternate" hreflang="en-DE"
        href="http://example.com/de/about-us" />
```

```
</head>
```

Let me point out a few potential gotchas:

1. The value of the hreflang attribute can contain either one or two components, with the second component separated from the first with a hyphen. The first entry is always the language (in ISO 639-1 format), and the (optional) second component is the geographic region (in ISO 3166-1 Alpha 2 format).

Running through the example from top to bottom, we have:

- de: German language content, independent of geographic region
- en-GB: English language content, for visitors in Great Britain
- en-DE: English language content, for visitors in Germany

Now time for the pop quiz: Does hreflang="be" specify content for Belgium? The answer is "no"! That's because the hreflang format places the language first and "be" targets the Belarusian language. To target French speakers in Belgium, you would need "fr-be".

2. Every page must identify all internationalised variations, *including itself*. As you'll notice above, the second entry pointed back to the URL that was currently being accessed.

3. If the hreflang entries in page A link to a regional variation in page B, then the hreflang entries in page B must also link back to page A. Google calls these "confirmation links" and warns that annotations that are missing these pointers back may be ignored. The upshot of this is that the exact same block of hreflang links you saw in the above example (for the URL "example.com/en/about-us") must also appear on the URLs "example.com/de/ueber-uns" and "example.com/de/about-us". Web applications that have built-in fragment caching systems could cache and share these snippets across the related pages instead of regenerating them afresh every time, thereby bolstering

www.dbooks.org

performance.

Viewing your internationalised Google search results

It's possible to check whether your hreflang efforts were successfully understood by Google through their Search Console. After logging in, navigate to Search Traffic > International Targeting, and you should see a readout if everything is in order.

Because Google Search remembers and detects your region and preferred language, it can be difficult to see how your website presents itself in searches originating from other regions or done through other languages.

Luckily, there's a helpful tool for doing this: Google's Ad Preview.[4] Ostensibly designed for viewing *adverts* in faraway lands, this tool also displays regionalised organic search results—which is very handy for evaluating internationalisation measures.

[4]http://adwords.google.com/d/AdPreview/

11 | Copywriting

Copywriting is mostly about persuading potential customers to buy from you. It is best conceptualised as salesmanship in print. Your potential customers have (potentially serious) problems to solve, and the job of your copy is to present your product as a credible solution. This is done partially through rational appeals to these potential customers' self-interest, and, if we copywriters are to be honest about what we do, partially through exploiting psychological tendencies. The latter, which is sometimes known as "the dark arts", may not sit well with the most upright of advertisers, and their offence is understandable. In defence of copywriters who dabble in the dark arts, their behaviour becomes less reprehensible if their product happens to be the best on the market. Here, there may be almost a moral imperative for them to ensure that the customer chooses their offering over some inferior substitute, especially when that alternative is so bad as to be dangerous.

In spite of its name, copywriting doesn't just refer to the *text* in your sales pages—although that's obviously a hugely important component. It also extends to your choice of images, videos, logo, colours, webpage design, and even sound effects. And copywriting is everywhere—not just on the pages of your website, but also in your advertisements, your brochures, your about page, your auto-appended email footers, your website terms & conditions, and your company's automated answering machine.

I'd like to share some statistics about the relative behaviour of people who buy something on my website versus those of people who don't. The point of this will be to demonstrate just how *much* of your copy the average customer will devour before buying—and just how *little* an uninterested person glances at before leaving. The average customer spends 25 minutes and 52 seconds on the website before buying, reading 11 pages. At an average read-

www.dbooks.org

ing speed of 300 words per minute, this amounts to about 7500 words. By contrast, the average non-customer spends 7 seconds viewing just a single page. This only amounts to about 35 words (i.e., just a few headlines). This comparison should show that your headlines are critical in getting visitors to stick around and read further. But once they're hooked, they'll read nearly everything you've written.

Emphasise Benefits Over Features

The engineers and designers of products have a blind spot: They are predisposed to thinking of their creations in terms of the features they worked hard to implement, rather than the pain-points that motivated these features' inclusion in the first place. They speak solely about their products' specifications, leaving out explanations about how those specs will benefit the customer. But in failing to emphasise these benefits, they hamper their sales.

Let's continue by analysing a few examples. Amazon's subscription delivery service Amazon Prime translates their feature "Next-day delivery" into benefit-speak like "Need a last-minute gift? Can't get out of the house? Realize you forgot to pick up something?" Switching our focus from the bookseller to the books being sold, is it a coincidence that some of the bestselling self-help books have embarrassingly direct book titles? *Think and Grow Rich*. *How to Win Friends and Influence People*. *The Seven Habits of Highly Effective People*. No one would be caught dead reading a copy of something with such a name in public, yet that doesn't stop these books from selling astonishing numbers of copies.

Writing copy in terms of benefits can feel a bit weird because it involves brainstorming all the possible good things that can come to customers who use your product. Often, this amounts to stating what is, from your perspective, totally obvious, and proceeding to spell out, step-by-step, how feature X leads to benefit Y. Those most fluent in benefit-speak are able to think up *second order* benefits (i.e., what are the benefits to having the benefit?). Sup-

pose we are selling a protein bar. Its main feature is that it contains 40g of balanced amino acids. The basic benefit to the typical buyer is that the bar will help them gain muscle. The second-order benefit, then, is that they will look better with their shirt off. And if we take this a step further, the third-order benefit would be feeling more manly and confident. Later-order benefits tend increasingly towards emotional ends—the kind of deep-seated wants and urges that motivate people.

What are some examples of these deep-seated human desires anyway? Sex (diet plans), companionship (UK pub culture), health (gyms and spas), freedom from physical pain (acupuncture), superiority & social status (expensive handbags), physical comfort (air conditioning units), the enjoyment of sensual pleasures (delicious tasting meals), the facilitation of laziness (food delivery), the accumulation of wealth (stock market brokerage), and self-actualisation (meditation retreats). If you'd like to look further into human desires, I suggest you visit the work of the anthropologist Donald Brown, who, in his book *Human Universals*, describes how every culture he studied shares some universal traits.

Advertisers appealing to second- or third-order benefits cannot always come out and make these claims explicitly; this may come across as cheesy or unconvincing. Instead, these advertisers communicate their messages through subtler means, such as by association or symbolism. We will examine these mechanisms elsewhere in this chapter.

Above, I talked about selling to consumers. With regards to selling to *businesses*, you needn't go so far into benefits-speak; business buyers already know their own problems, and have a rough idea about how the solution will look. When marketing to them, you are better off focusing on how you can save them time, money, and risk.

Not-for-profits with millions of pounds in their bank accounts are known to go bankrupt from time to time. How can this be possible? The problem is that whenever a not-for-profit receives

www.dbooks.org

donations, it's often exclusively for one particular purpose (e.g., for the purchase of art or the upkeep of the premises). It's illegal for the not-for-profit to take money from one of these accounts and spend it in another (e.g., for the paying of staff salaries). In much the same way, your customers (private individuals and businesses alike) have a set of mental checking accounts, each intended only for buying one category of thing, and given its own budget. Two hundred pounds might be someone's upper limit for replacing a work-related dress suit, whereas that same person wouldn't think twice about allocating a budget of a thousand to making himself attractive to potential partners. The salesperson who frames the suit purchase in terms of attractiveness may sell a more expensive suit, showing that the choice of *which* benefit to emphasise affects end results.

Overall, writing about benefits is a more inventive copywriting process than describing features. The only qualification for feature-speak is that you're good with a measuring tape and can be trusted to accurately pass on information you get from elsewhere. Benefit-speak demands that you get into your customer's skin—that you play the armchair psychologist.

Here are a few more examples (from a fictional smartphone manufacturer) that help delineate the difference between features and benefits:

Feature	Benefit
902.11ac wireless	Access the internet when on holidays or doing business abroad—even in places that older phones didn't support.
B8 processor	Your battery will last significantly longer without needing to charge, meaning heavy-use days or unexpected sleepovers are less of a hassle.
Super-Duper finger ID technology	You don't need to remember PIN codes anymore…no more getting locked out of your phone and having to rummage for your PUK code.

Feature	Benefit
Bronze graphics technology	This phone's video games will match current generation games consoles in 3D graphics realism. You'll be able to play blockbuster titles like X.

Before concluding this section, I'd like to share something written by copywriter Clayton Makepeace. He sarcastically illuminates the difference between features and benefits in the context of teaching readers how to market a fictitious hormone-balancing medical product:[1]

> *Have you ever been awakened in the middle of the night...sat bolt upright in bed...slapped yourself on the forehead and exclaimed, "Holy moley—I gotta get off of the hormone roller coaster?"*
>
> *No? Me neither!*
>
> *Have you ever found yourself feeling eager to PAY for a product that would do any of those things for you?*
>
> *Nope? Join the club*
>
> ...
>
> *Our "hormone balancing" prospects want to stop having hot flashes and mood swings and stop losing their libidos.*
>
> *Why? Well, for one thing, because hot flashes and mood swings are irritating—even miserable. And for another—drilling down even deeper—because all of these things threaten the intimacy and security of their primary relationships. Nobody wants to be a hormone hermit!*
>
> ...

[1] http://www.makepeacetotalpackage.com/archives/make-your-products-benefits-sparkle/

www.dbooks.org

> *Nobody really wants to balance their blood sugar levels. But anyone in his or her right mind DOES want to avoid the misery of blindness ... cold, numb, painful limbs ... amputation ... and premature death that go along with diabetes.*

Even if we understand the basic difference between benefits and features, our well-intentioned attempts at benefit-speak can nevertheless mutate into feature-speak should we happen to define the problems our products solves at a level too decoupled from deep human desires. The extent to which this happens may depend on the audience you are marketing to. For example, a diet pill advertised as "increasing calories burned by 20%" would certainly be benefit-talk to someone who counts calories and is comfortable with math. But this same message may only be feature-talk to the mass of customers who are less than compelled by the promises of percentages. These customers would prefer a pill that unashamedly but clearly promises to help them "lose a stone in two months". All in all, it's better to leave the mechanism of action to the small-print, and instead focus on the WIIFM (What's In It For Me). Sophisticated consumers will simply shrug off your bold marketing claims as fluff and refer directly to the ingredients or to third-party reviews. But by brazenly advertising the ultimate benefits of using your product, you'll win over the multitude that can't be bothered with the details.

Convincingly Demonstrate Credibility

The sales process—be that in text or in person—begins with you getting your message in front of a prospective customer. In the next stage, you demonstrate how you plan to solve that person's problems. But before they'll buy, you need to convince them that your claims are credible, and that you really *can* help them with their problems.

In this section, we'll be focusing on that final hurdle before the sale—that of demonstrating credibility.

Have you ever visited one of those single-page sales letter-type websites that scroll on forever and ever and promise to reveal to

you the "secret to making money online through affiliate sales"? Where the author reels off hundreds of advantages in gaudy red and black fonts set in wildly fluctuating sizes and scripts, where ubiquitously glowing testimonials are submitted by people with suspiciously similar writing voices to that of the original author, and where the price is perennially "reduced" from £499.99 to £19.99, payable via Western Union Transfer? These websites lean into your ear and scream "SCAM".

And here we meet the first rule of credibility: Avoid practices, memes, and signals used by the less-than-credible denizens of the Net.

Suppose we turn this rule on its head and state it in the positive instead. This would prescribe that we incorporate into our copy the signals and signs used by those whose credibility is beyond questioning. Have you ever visited a major law firm's website? A typical sounding business name might be "Brown, Goldstein, & Partner". I'd like to point out a few things of interest. Notice that the firm's name is composed entirely out of classy sounding surnames, one of which (Brown) is also the name of an elite university, and another of which is suggestive of a precious metal (Goldstein). Notice the ampersand stuck in the middle, instead of the more normal "and". Notice the word "Partner", appearing in singular, no less! All of these factors trigger subconscious associations of trustworthiness in the minds of whoever encounters the name. In short, "Brown, Goldstein & Partner" sounds like the kind of name a legitimate law firm would have, whereas something like "John-Joe, Jimmy, and Their Buddies" definitely doesn't.

The copy on the fancy law firm's website also sends plenty of signals that help contribute to the sense of legitimacy we feel when encountering their brand. These include imagery of luxurious offices and of their lawyers in expensive suits, enterprisey vocabulary such as "client", "firm", and "solutions", and a web design that is uncluttered and every bit as conservative as the likely politics of the firm's lawyers.

www.dbooks.org

I mention all this not because I want you to copy a law firm's website for your business—this style of branding would hardly be fitting for a downtown nightclub or a kid's pizza place. Rather, I point this out because I want you to start paying attention to your feelings of trust and distrust as you browse the web. Notice the little things, the subtle signals that these websites send, like the tiny padlock above the mention of SSL, the word "secure" printed above photos of MasterCard, or the pictures of the business founders smiling at you from the about page.

So far, we've discussed getting your signalling in order. The second way you'll win credibility is through the social proof of reviews, customer testimonials, association with trusted institutions or people, affiliation with official dealerships[2], or showing off a flurry of human activity on your website pages. This last one is the online continuation of an age old practice wherein older sitcoms dubbed over their audio with canned laughter tracks, nightclubs engineered long and artificial queues, and restaurateurs sat their first customers of the night near the window for passers-by to see.

Mathematicians might prove their theories by formal processes like "proof by contradiction", but the rest of us rely on "proof by the Joneses" (i.e., the law of following the herd). Part of your job as a marketer, then, is to gather and solicit positive words about your product and then highlight these in your marketing materials. A long-standing favourite among webmasters is what I call the brand name-drop, where the webmaster name-drops (with logos) a bunch of major companies and publications that supposedly have positive things to say about the product. I say "supposedly" because, while most of these brand name-drops are legitimate, sometimes they can be surprisingly tenuous and shallow. For instance, I saw a website earlier today whose home page claimed "as featured in MAJOR NATIONAL NEWSPAPER X". No link to their mention was provided, and when I researched further I found that the website was merely mentioned along with

[2]E.g. Express Watches, a watch retailer in the UK, more than doubled their conversion rate by adding a prominent image that read "Seikko Authorized Dealer Site" - https://vwo.com/blog/increase-in-sales/

20 other start-ups in that paper's roundup article about local companies who attended XYZ event.

Popular too is the testimonial, a marketing device whose magic lies in its ability to voice claims so ridiculously glowing or promises so overly grandiose that no business could ever legitimately make them about themselves (e.g., "Changed my life", "A thousand times better than the rubbish by competitor X", etc.)

Testimonials need to be believable to work, so it's important to accompany them with identifying information about the testimonial writer (e.g., their photo, full name, workplace, and social media profile link). Believability reaches its fullest potential if the person writing the testimonial is already known, respected, or even revered by your potential customers.

Dubiosity is, unsurprisingly, rampant in the testimonial department. Besides outright faking, there is also exploitation of tenuous links and half-truths. For example, I know a marketer who presents positive words said by their golf buddy at Goldman Sachs as being the opinion of the Goldman Sachs institution as a whole.

Let's talk about website activity. A brand new web app without any users or content is the technological equivalent of a deserted amusement park in a depopulated post-apocalyptic world. In both, the lack of human movement and activity is eerie and supremely off-putting. The best way to prevent your website from giving off this sort of ghost town impression is by programming it to show user activity. This activity will be real if you've been around for long enough, but it might also be artificially generated if you've just launched and don't mind dabbling in the dark arts.

The devices used to demonstrate activity come in various forms, the most common being rapidly updated activity feeds, comment feeds, recently listed product feeds, etc. On top of this, aggregate statistics that display your particular business impressively can also be helpful (e.g., displaying the amount of money

www.dbooks.org

paid to platform partners, the average number of dollars saved, the number of website templates for sale in your system, the number of to-dos completed, the number of seconds saved, and so on and so forth).

The next major point about building trust is correspondence to and accountability in the offline world. The website of your typical law firm will transparently and prominently identify themselves, their whereabouts, and how to reach (or sue) them. The law firm's address and phone number are displayed in the website footer, and a quick gander at Google Maps Streetview confirms that an office of the same name exists at that same address. The firm's website lists their company registration number and sales tax number—both of which correspond to the data at the company registrations office. They flaunt the fact that they were incorporated 60 years ago, signalling that they must be competent and trustworthy enough to stay in business for such a long stretch of time.

So far, the above shows that the institution has made itself accountable. But the law firm's website goes further and also communicates that the individuals *constituting* the institution are also worthy of your trust. The website contains photos of their executive team, and upon googling them you'll find the normal spattering of social media profile stubs and photos from charity balls, weddings, and half-marathons. In short, everything is very real.

The underlying message in all of this is that the law firm has nothing to hide. Compare all this to the red-and-black sales letter website, where there is no identifying information, no office address, no phone number, no team photos, and no sign that the authors exist in real life.

There's a tension here for the lifestyle entrepreneur, that breed of modern business owner who wishes to live without the ties of fixed office or schedule. This digital nomad doesn't want customers calling in on them in person or on the phone. They don't want anyone interfering unexpectedly with their time. One option for them is to just accept that their desired lifestyle comes

with inevitable trade-offs in credibility signalling. Another option, residing in the grey arts, is to resort to hacks and other clever devices. This entrepreneur will leave a telephone number on their website, but this line only connects to a Skype answering machine, enabling the entrepreneur to return the call whenever it suits them. Similar hacks could be done with addresses (e.g., PO boxes, etc.)

The fourth big way to build credibility is to signal confidence in a costly manner—essentially, to put your money where your mouth is. Consider, for example, lifetime warranties and no-questions-asked money-back guarantees. It doesn't take a game theorist to realise that any company confident enough to make these promises must believe that their product is top notch—otherwise, these guarantees could end up being used so extensively as to put them out of business. In effect, the very *existence* of a guarantee signals that the business is confident about the quality of their products, and this increases sales. Anecdotally, some companies even report *increases* in profits following the introduction of money-back guarantees. For them, the improved conversion rate built on their newly boosted credibility more than offsets the losses through refunds. (That said, money-back guarantees might not be such a good idea in markets filled with serial refund abusers.)

My very last general point about demonstrating credibility is that timing matters—it's especially important to emphasise your trustworthiness at that moment before your potential customer will purchase. That's why so many websites have what are called "click triggers"—start ratings, micro-testimonials, or assurances about the money-back guarantee—placed next to their buy buttons.

Give Reasons to Act Now as Opposed to Later

Which one of this pair is busier? The surgeon staffing the emergency room of a chronically understaffed hospital or the social media community manager? From my standpoint, I would wa-

www.dbooks.org

ger that the surgeon is busier, yet both she and the community manager, when asked about their recent work, will report that they have been "busier than ever". I bring this up to show that everyone considers themselves busy today—even if, comparatively speaking, they are not.

If you combine the modern constancy in self-reported feelings of busyness and the inevitability of distraction in today's attention-scarce world, this creates a problem for online salespeople: Many leads, even if they adore the product and are on the verge of buying it, may nevertheless drop out of the sales funnel. All it takes to lose them is their receiving a text message from their crush or getting a Facebook notification or having the thought that your longish sign-up form would be easier to fill in when they're at home later. Once you lose their attention, it's unlikely you'll ever get it back.

The introduction of scarcity counters this kind of fickleness by putting on the pressure to act *now* as opposed to *later*. Fearing that the scarcity will take away the opportunity completely, the lead ignores all other distractions until they complete the sale.

Broadly speaking, scarcity comes in two varieties: unit scarcity (e.g., Amazon tells you there are "only six left in stock") and time scarcity (e.g., a clothes retailer tells you that their "sale ends in three hours").

Unit scarcity works partially on the basis of social proof. People see low availability and conclude that other people must be buying the product in droves. And if it's good enough for all them, it can't be a bad choice.

Unit scarcity also works because people have a psychological aversion to losing freedoms—even if the freedom is merely the ability to own something they didn't even want in the first place. Threatened with this kind of loss, people react, almost automatically, by striving to possess.

Scarcity tactics has been around for a long, long time. Cialdini,

in his must-read book *Influence,* provides some examples from the pre-internet age: Real estate agents pretend a doctor and his wife from out of town are interested in the property you want to buy, or photographers at baptisms claim that "stocking limitations" force them to dispose of unsold photos within 24 hours.

The more plausible the scarcity is, the better. Real estate agents prove they aren't bluffing about demand for a property by scheduling a crowd of other interested parties to view the apartment at the same time as you. Or, if they don't want to go that far, they will, at the very least, ensure that some other suitor "accidentally" crosses your path on the way out of the apartment viewing.

If you're running a fully digitised software business, there aren't so many opportunities for natural scarcity. Dabblers in the dark arts will thus seek to manufacture scarcity in the most plausible manner they can. On the low-end of the plausibility spectrum are marketers who limit digital downloads of their course to X customers per month in order to "stop the information from getting into so many hands that it loses its powers".[3] More plausible are web retailers with one-day-only 50% discounts, or those who introduce scarcity via coupons which are "limited to the first 50 claimants". Another plausible road to scarcity is the "beta launch", where the marketer only accepts 100 customers to their pre-launch product, on the grounds that their support team won't be able to handle any more than that at this early stage. Just as plausible is shipping scarcity: Some retailers offer free shipping if you order within X hours, and to make sure you don't forget this, they conveniently place a ticking countdown timer next to their buy button.[4]

My experiments with using scarcity in my marketing had the effect of improving conversion rates in author recruitment from

[3]For instance, see author and pick-up artist Neil Strauss, who launched the world's most terrifyingly named dating course, Annihilation Method, with an artificial limit of 375 copies. Although many people complained about this artificial limit, reportedly he made serious money.

[4]https://www.whichtestwon.com/case-study/the-count-is-on/

www.dbooks.org

about 8% to about 35%—a phenomenal jump that dwarfed any other improvement I ever made to this process. My author recruitment flow works as such: Applicants apply to publish their study notes with my company, and we review these applications, accepting those who fit our criteria. Around 2011, we decided to start radically limiting the number of authors we accepted, and we communicated as much in the copy on our author application page. The explanatory reasoning we gave was that 1) too many authors in our system causes too many products to be generated, and too many products causes potential *buyers* to be overwhelmed and paralyzed by the extensive choice;[5] 2) each additional author creates editing and formatting work for our team to deal with, and we need to limit these costs; 3) we don't have the manpower to consider author applications all year, so we have to choose a cut-off point (the start of the academic year, September 1), at which time we won't consider any later applications.

We also use some scarcity within the later stages of our application process. In the contract we send to potential authors, we communicate that they have three days to accept our terms, otherwise we cannot guarantee that their place won't be given to another author. We justify this because we are under time constraints to finalise our content before the academic year starts.

Prefer Concrete, Visual, Specific Copy

Coming from a programmer background and mind-set, my tendency is towards seeing the world in abstractions. Ask me what happened in any Marvel or DC Comics action movie—be that *Batman*, *Spiderman*, or *The X-Men*—and I'd say, "A hero saves a city from danger". I mercilessly scrape away the details that differentiate each movie, leaving in my wake a lifeless and bland generality.

This is exactly what not to do with copywriting. Why?

First, by writing abstractly, we fail to differentiate our products

[5]https://www.wikiwand.com/en/The_Paradox_of_Choice

from competing offerings. If all three action movies are about heroes saving the day, then why does it matter which one I watch? Similarly, if 10 other websites are promising "the best ebook on vegan diets" just like we are, then a customer has no reason to choose ours above the others. To stand out, we need to provide memorable and compelling additional details.

Second, the human mind remembers specific, unique, or weird details more easily than dry abstractions. Mnemonic technique, as espoused in the bestselling book *Moonwalking with Einstein*, emphasises this very point. This thinking predicts that a customer's memory circuits would remain under-stimulated with a title like *The Best Ebook on Vegan Diets*, but would be brimming with activity with a title like *What the Dietary Habits of Herbivorous Tibetan Bears Can Teach Vegans About Healthy Eating*.

In fact, when I read the title *The Best Ebook on Vegan Diet,* the first thing I see inside my mind's eye is...nothing! This is because my mental machinery defaults to laziness. If it can avoid constructing detailed images, then it bloody well will. Without something concrete being etched into my memory circuits by the book title, I will be prone to quickly forgetting it. Indeed, even if I consciously *force* myself to imagine something in response to this bland title, my mind, being lazy, will likely choose from its currency of stock symbols, images, and ideas, reusing images from my other memories rather than inventing afresh. As such, the evocative novelty of anything I conjure up will be minimal, for I will only be seeing the clip-art of my imagination—hardly likely to be memorable.

Third, generality and abstraction is easy to fake, and this could mean that the customer will suspect being lied to—even if the message were in fact totally honest. Think about bland advertising claims such as "Supreme quality", "Best in California", or "Better ingredients". Totally unconvincing, right? Everyone makes the same flaccid claims, and it's as if the cheapness of the marketer's imagination seems to predict a cheapness in the quality of the product. On the other hand, think about a specific, non-

www.dbooks.org

fluffy claim like "We use (maximum) three-day-old hand-picked Brazilian acai berries in our homemade smoothies, according to the CTC Healthy Practice juice-making guidelines". That sounds a whole lot more wholesome and believable.

Specific, visual language has yet another advantage compared to summary truths of the same ideas: Our emotional triggers are far more responsive to vivid stories and anecdotes that an unadorned statement of the same facts. We all "know" that millions of children starving is a dreadful state of affairs, yet we only feel motivated to pull out our credit cards to donate after seeing a televised charity advertisement that depicts a forlorn child staring at the camera. When we go to copy-write—which could well be on behalf of such a charity—we too should consciously pull these emotional heartstrings. This could be as simple as a case study of how "this one guy" (who happens to be really similar to the customer) found, partook in, and was positively changed through using our service.

Concreteness mostly applies to text, but also applies to other aspects of copy. ExactTarget, who in 2013 were hosting a conference for marketers, increased their registration rate by 40.18% when they replaced a photo of a bland, contextless laptop with a photo of a crowd of attendees in a crowded conference hall.[6] Their original photo just didn't create the right visuals to fire up their potential attendees. Another way to spruce up vividness is by creating a brand mascot. Mailchimp have their cute little monkey, just as Kellogg's have Tony the Tiger. The existence of these mascots makes these brands easier to conceptualise (and remember), compared to brands without a memorable face.

I'll conclude with a rough rule of thumb for measuring the extent of your copy's specificity: Check that your text contains perceivable concrete nouns ("bears", "monthly energy bills", "cups of coffee"), numerical quantities ("12-second tooth-brushing session", "prices reduced by 35% since January", "our profit margin is only 5%", "established in 1994", "27 cheesy sandwich ideas you HAVE

[6]https://vwo.com/blog/choosing-images-that-improve-conversions/

to try before you die"[7]), proper nouns ("Golden Gate Bridge", "Arnold Schwarzenegger"), and photos of activity.

Ensure Scannability, Readability, and Responsiveness

When composing sales copy, the part of you that appreciates slow-paced, nuanced, paragraphed prose must surrender the writer's seat to the part that is into frenzied lists, one-sentence paragraphs, and zealous use of HTML formatting.

The reader of a large, unformatted block of copy cannot quickly extract meaning. From a sales perspective, this means that they cannot quickly learn how your product will benefit them. They may deplete their patience after a few seconds and decide to give your website a skip, only to shack up with a competitor who took the time to present their pitch in a more readily scannable fashion.

To stand a chance of communicating with the modern web surfer who will make a snap judgement in the first five or so seconds of landing on your page, you ought to present your copy using numbered lists, copious descriptive headlines and taglines, generous bolding/highlighting/underlining, and liberal use of white space and document structure. On top of this, you should be aware that some parts of the page are especially salient, and you should therefore strive to place your most important marketing messages in these locations (e.g., in your headlines, image captions, button texts, or within any text boxes designed in ways that contrast visually with the rest of the page).

Due to systematic biases in how humans make judgements, the cognitive ease with which we can see, read, and absorb meaning from text can spill over and transform into positive feelings about the *content* of the text.[8] In other words, the very fact that a message is readily understood nudges it ever so slightly towards

[7]Headlines with numbers in them "resonated" twice as well as normal ones without numbers, according to a Conductor.com survey. Regrettably, their data was not based on actual click data. https://moz.com/blog/5-data-insights-into-the-headlines-readers-click

[8]For a good overview, read *Thinking, Fast and Slow* by Daniel Kahneman.

www.dbooks.org

being believed and embraced.

This means that we can enhance our copy's persuasiveness by making it easier to comprehend—for example, by writing in larger font sizes, colouring fonts so that they contrast strongly with the background, limiting ourselves to familiar vocabulary, shortening sentences, painting analogies with familiar rather than distant concepts, writing concretely as opposed to abstractly, laying off acronyms and jargon, imposing structure on your text, and supplementing everything with an ungodly number of illustrative photos.

Of course, all the formatting and structuring in the world won't do your sales copy one bit of good if your design gets bungled on certain devices. Act-On Software used to send their mailing list subscribers newsletters with layouts that were so unresponsive that anyone on a mobile email client had to scroll left and right in order to read each and every sentence. When this company switched to a design that rendered responsively on mobile devices, they improved their conversion rates by a whopping 130%.[9]

Although responsiveness on mobile devices is the topic of the day, in the future there will no doubt be new devices and ways for people to experience the web. Marketers who are aware of and ready to optimise for these upcoming technologies will enjoy superior conversion rates.[10]

Signal with Indirect or Unworded Messages

Not all communication is explicitly wrapped up into words. In our personal lives, we constantly send implicit signals to those around us by means of the clothes we wear, the cars we drive, and the company we keep. An expensive Rolex communicates status and wealth, as does an address in an expensive part of town.

Some signals are sent as unintended by-products of our lifestyle choices and activities, whereas other signals are manufactured

[9]https://litmus.com/blog/responsive-design-ab-testing-leads-to-a-130-increase-in-clicks
[10]Google Analytics data will inform you of the exact breakdown.

with the explicit goal of influencing perceptions and impressions. A man riding a motorbike will be perceived by many as rebellious, and this is so regardless of whether this was his intention or whether his only reason for riding the motorbike was that he couldn't afford to drive a car.

Why signal? Mainly because it is more socially acceptable to communicate certain (usually self-congratulatory) things by implicit means than by explicit statements. You'd view a wealthy person as obnoxious and rude for bragging about their net worth, but you'd have no beef with them for inviting you to hang out at their Beverly Hills mansion—yet, in both cases, this wealthy person will have communicated to you their wealth.

Businesses use signals too when they want to communicate facts that might be inappropriate to mention directly. A prominent "we're hiring" link signals that the company is cash-flow positive and growing, as does a flashy office in the best part of town.

There is potential for dishonest signalling, a concept from the natural world whereby signals are manipulated to give a false impression. An example of this might be a non-stinging bug that mimics the yellow and black stripes of a wasp and relies on this to intimidate its predators. Back inside the office world, I know a young consultant who launched her services with an "about page" containing 15 employee photos, despite it only being her and her alone at the company. She wished to attract big clients and thought that the best way would be to pretend that she was a big deal. (Five years later: she now *legitimately* employs 50 people–obviously her tactic worked!)

Explicitly Point Out Qualities, Lest They Go Unnoticed

Communication is a brutally mechanical process. Despite the claims to the to the contrary online,[11] there is no such thing as telepathy. For communication to happen, our thoughts must be gathered, packaged into words understandable by others, and

[11] http://www.wikihow.com/Develop-Telepathy

www.dbooks.org

then delivered, without loss, to our listeners and readers.

Yet owing to a mental shortcoming known as the "curse of knowledge", we systematically underestimate the difficulty of communication and assume that others magically know what we know.[12]

Part of the challenge in writing copy is overcoming this curse. To do this, we need to emulate the mind of a complete outsider, someone who has not shared the same intellectual and experiential journey as us, who hasn't seen what our eyes have seen.

One popular trick that helps in this respect is to describe your offering in *much* greater detail that your natural tendency would suggest. This entails writing as if your customer were a total newbie to the market and your job were to educate them not just about your product, but also about what each of your features means, why they are significant, and so on. By way of example, imagine that you and a competitor are both offering an identical service, indistinguishable in every way and in all of its accompanying benefits. The only difference between you and them is that your company—and your company alone—*mentions* the following points in its copy:

- "We provide official tax deductible receipts for business persons (and students), allowing you to offset your purchase against your next tax bill."
- "We allow customers to re-download their ebooks for up to three years after purchase."
- "We answer all customer service emails in 24 hours."
- "We provide our ebooks in three formats—PDF, MOBI, and EPUB—ensuring you can read on any device."
- "We provide secure payments with industry-grade HTTPS security."

Let me reiterate that your competitor provides these exact same features and benefits; they merely fail to mention them in their sales copy, thinking them "obvious" or "unworthy of mention-

[12]https://www.wikiwand.com/en/Curse_of_knowledge

ing". The consumer who visits both your website and your competitor's will simply assume that your competitor doesn't offer the advantages that you do. The result: The customer chooses to buy from your company instead because it seems like you are giving them more for their money.

Some companies throw in freebies with their orders (e.g., "If you buy this mobile phone, you get this headset for free" or "If you buy this acoustic piano, we will deliver it free of charge"). It's important for these companies to mention *how much* their freebie would cost on the open market, otherwise it risks being undervalued by the potential customer.

As well as mentioning advantages and freebies, it is also important that a business remember to tell its customers what it went through to *produce* whatever it is offering.

Let's imagine you are an entrepreneur in the clean-tech space running a website that retails residential solar panels to private consumers. Before launching your business, you travelled the world and vetted 200 solar panel vendors, choosing the 20 that offered the very best panels in terms of efficiency and reliability. This background research and product filtering is obvious to you, since you did it, but, without any explicit mention of it in your sales copy, your customers will remain completely oblivious to your journey and the advantages it brings them as consumers. Indeed, they might even scoff mockingly at the seemingly meagre range of panels your business stocks when compared to the competitor that has no vetting process or standards to speak of.

There is no need to be shy about telling your customers, in detail, about the methodology you used to choose the best panels to stock. Tell them about how many doors you knocked on, how many reviews you trudged through, and how many spreadsheets full of figures you crunched. Share with them the criteria you employed in your decision-making formulae so that they can see for themselves. Reveal to them the lofty qualifications of the engineers on your evaluation team, so that the customer becomes aware of the expertise backing your service.

www.dbooks.org

For those of you working in IT, the physical aspect of the production process—over-caffeinated coders sitting at computers in a humid office—may not be so glamorous. But there is still plenty of opportunity to be conceptually dramatic. For instance, why not take a leaf out of Google's book and *name* and describe the great things about your algorithm, as they did with *PageRank* and the *Penguin* update?

One last thing: It's also important to inform the customer about what went into your product *at just the right time*. If you only reveal these compelling details way down in the depths of your about page, you cannot expect the average customer to read them. Instead, you would be better off briefly describing your production process on your landing page, along with a link that leads to a meatier description for those readers interested in knowing more.

Compare Yourself to the Competition

Presenting a believable, professional, and balanced comparison helps win you the customer's trust by showing that your company has nothing to hide, and that you're open and transparent and not afraid of the competition. It's also possible that, upon reading your comparison, the customer will conclude that all the research was already satisfactorily done for them so they need not bother checking the competitor's website. These people will finish the sale right then with you instead of potentially getting distracted while they research competitors.

There is also an SEO advantage to comparisons: Some of the most popular search queries are in the form "X alternative", "X vs Y", or "X comparison". SEO-optimised pages on your website that directly answer these questions are sure to capture a slice of this highly relevant traffic.

A source of morally questionable leverage lies in the marketer's choice of comparator. This moral issue is most often encountered in pharmacological research, where big pharmaceutical companies compare their latest anti-depressant pills against a

placebo instead of against the most effective cure currently on the market (be that another pill or alternatives like exercise or therapy). Their choice of comparator, while proving that their medicine works, doesn't prove that it's as good as the other offerings on the market. Returning to the world of website copy, it's up to the marketer whether they want to *only* compare themselves to the big stagnant market-leaders that everyone knows and hates, or whether they want to also compare themselves to the impressive upstarts.

Assuming you have the appetite for comparative advertising (see the warnings below), the easiest way to get started is to create a table or chart comparing your company to your competitors. This format is useful because customers can read off, at a cognitively easy glance, the factors that elevate you above everyone else serving the market.

Comparing yourself with competitors, while effective, opens you up to the risk of getting sued. When making claims about another company, you need to be absolutely certain that what you say is true and substantiated (i.e., download and keep evidence). It may not even be enough for your claims to have been true at the time of writing—there has been litigation over claims that simply fell out of data. (Those working in the software industry can at least point to a version number, thereby protecting themselves.) Please don't interpret any of this as legal advice—ask a lawyer.

Furthermore, when mentioning a competitor, you are required to include a disclaimer that you are unaffiliated with you them, just in case a customer sees your competitor's branding and mistakenly believes your website to be theirs. There are even more rules about this, so do yourself a favour and research them thoroughly before writing your copy.

Even if you don't have the risk appetite to directly compare yourself with your competitors, there are more subtle means to get a good word in on the back of their brands. One method is to mention a competitor in a neutral or positive light, simply to pick up the SEO advantage. I know of a website that did this and out-

www.dbooks.org

ranked their competitor *for the competitor's own brand name*! Another method for picking up a competitor's SEO juice is through your on-site testimonials. These testimonials can bash competitors as much as their writers wish, as long as you, the website owner, disclaim responsibility for and affiliation with their opinions. (But again, check with a lawyer.)

Anticipate and Counter Objections

Years ago, I remember seeing a little-known book on Amazon that had the scammiest looking cover I'd ever seen, yet had been rated with nothing but five-star reviews. Even though I was suspicious, I ended up being swayed to buy the book after reading the following (also five-star) review:

> *"I was sceptical when I noticed that the book had only five-star ratings. I wondered if the reviews were fake. By the time I was halfway through the second chapter, I realised that this book delivered."*

Like many people, I am always unsure whether online reviews by random internet users can be trusted, especially given how easy they are to fake. But the review I just quoted was special because it faced this objection of mine head-on and dispelled it with two succinct sentences. (Or at least, largely dispelled it, since it's possible that a *really* sneaky book author could have also faked the above so as to make his other fake reviews look that much more believable. In this case, my scepticism was unfounded. The book in question, *Ca$hvertising*, was superb.)

In classical sales, you are standing next to the customer in conversation. It's possible to read their reactions, address their misgivings, disabuse them of unfounded fears, and generally soothe them through the process. With copywriting, there is a physical distance between you and the customer, so the conversation necessarily begins and ends with your copy. The only chance the copywriter has to address the customer's fears is through their text—and this is why good copywriters don't shy away from talking about the most common objections.

The downside of raising objections in your copywriting is—surprise surprise—that you risk bringing up a disadvantage that the customer wouldn't have thought of, and this could scare them away. But perhaps this isn't as much of a downside as one might think. If the customer had not considered a major downside, then they might suffer buyer's remorse and demand their money back or leave negative reviews online. Addressing objections, then, is also about setting appropriate expectations.

Bonus Points: Mention a negative, then a positive.

In cases where the objection you are raising points to a drawback, you can lessen the impact by mentioning how this drawback is related to some advantage. For instance, think of the undersized apartment that gets advertised as "small but cosy" or the unadventurous meal described as "simple but wholesome".[13]

Associate Yourself with Desirable Connotations

When a sports star wears Nike trainers on the field, the message communicated to onlookers, if it were to be written out explicitly, might go something like this: "I wear Nike and, as you can see, I am a famous, edgy, sexually desirable, and super successful sports star. If you were to wear Nike clothing, then you too could be just like me!" The "logic" of these messages is little more than a blurring of correlation and causation. Indeed, during my more cynical hours (and there are many), I imagine that a handbook on logical fallacies could also serve as a handbook on advertising techniques.

The idea that advertising by association could possibly work at all can be baffling, yet advertisers of clothing, alcoholic beverages, and even university educations have been slathering their advertisements with associative symbolism for decades. Fashion retailer Abercrombie & Fitch built their brand on depicting beautiful, young, wealthy Americans enjoying the summer at their lake cabins. Alcohol distributor Smirnoff creates ads showing at-

[13] This point comes from *Yes!: 50 Scientifically Proven Ways to Be Persuasive* by Goldstein, Martin, and Cialdini.

www.dbooks.org

tractive finance professionals clinking glasses in trendy Manhattan nightclubs, and there is more than a whiff of sexual adventure in the air.

Given that large, economically rational agents like these employ and have continued to employ these techniques for years, there is plenty of reason to believe advertising by association is effective. I would hazard a guess that these advertisements work by activating automatic associative machinery in our brains. On the rational and conscious level, we "know" that depicting someone drinking a Corona beer on a beach is hardly an argument for Corona being the most refreshing beer to drink there. Yet, like a virus infecting a computer, the association between Corona and beaches bypasses our logical firewall and leaves its mark in our thoughts. Because Corona's advertisements constantly show us images of their beer next to images of beaches, our unconscious brains simply note that these two concepts go together, without asking why that should be so. And because this connection by co-occurrence goes both ways, the next time we are sitting on a beach, we cannot *not* think about having a Corona.

This isn't the only theory about why advertising by association works. Kevin Simler, in his fascinating piece *Ads Don't Work That Way*, argues that advertisements change the landscape of culture meanings, and we consciously align ourselves with the meanings that benefit us most:[14]

> *Cultural imprinting is the mechanism whereby an ad, rather than trying to change our minds individually, instead changes the landscape of cultural meanings—which in turn changes how we are perceived by others when we use a product. Whether you drink Corona or Heineken or Budweiser "says" something about you. But you aren't in control of that message; it just sits there, out in the world, having been imprinted on the broader culture by an ad campaign. It's then up to you to decide whether you want to align yourself with it. Do you want to be seen as a "chill" person? Then bring Corona*

[14] http://www.meltingasphalt.com/ads-dont-work-that-way/

> *to a party. Or maybe "chill" doesn't work for you, based on your individual social niche—and if so, your winning (EV-maximizing) move is to look for some other beer. But that's ok, because a successful ad campaign doesn't need to work on everybody. It just needs to work on net—by turning "Product X" into a more winning option, for a broader demographic, than it was before the campaign.*

Theory aside, let's get down to practice. How does one engineer one's own associations? There are overt means, such as photos and illustrations, endorsements, or choice of examples in your copy. The background is at least as important as the foreground here: Think of the previously regional politician posing for a photo in front of a map of the USA, symbolising that they have "gone national". Then there are subtle means, such as colour schemes, typography, vocabulary, and writing style. Petrol companies greenwash their corporate branding in leafy shades so as to suggest eco-friendliness. Protein powder manufacturers use aggressive, meaty, size 80 fonts to impress on their market that their powder leads to similarly oversized muscles. Academic and legal writers wrap their prose in dense abstractions, cold solemnity, and obscure vocabulary to suggest their intellectual might. Tabloid and blog writers buck that trend and write in a loose, conversational tone, throwing in puns and smileys and even the occasional profanity.

There is also a lot of subtle associative power in your brand and product names. Naming consultants make earth-shattering sums of money for coining words that evoke desired subconscious associations. Think about the name of the management consultancy "Accenture". The first syllable, "acc", reminds me of the word "accelerate"; and the last, "ture", recalls "venture". If this were accidental choice, I would be surprised. I've got another naming example: A friend invented an algorithm to figure out the cheapest renewable energy components for personal homes. If they christened the algorithm "Deep Green", then they would associate themselves both with environmental friendliness ("Green") and with the "Deep Blue" algorithm, the first pro-

www.dbooks.org

gram to beat a reigning human chess champion.

There is another breed of advertising by association that appeals to people's reason and rationality, as opposed to exploiting their unconscious machinery or need for social posturing. This rational appeal happens when a brand associates itself with a trusted expert or authority. For instance, think of the toothbrush manufacturer who claims that dentists use their brushes in caring for their own teeth. This message is compelling because a layperson has no way of evaluating toothbrushes, and so they must rely on expert advice in their decisions. And who better to follow than a dentist? What's more, the fact that the dentists *uses* the toothbrush themselves is even more convincing than their merely recommending it. After all, actions speak louder than words—even if said actions were not directly observed but instead were only recounted.

Repeat Yourself

The mere repetition of a claim increases the likelihood that those reading it will believe it to be true.[15] This last little fact is more than a little sad for humanity, but is great news for the advertiser prepared to dabble in the dark arts.

Selling by using repetition isn't as simple as copying and pasting the message to different parts of the page. If you do this, you'll come across as lazy, sloppy, and way too sleazy—not to mention boring. The art of repetition lies in tastefully drawing out your copy so as to allow the repeats to show themselves in more natural and entertaining settings. This is done mostly through the use of pleasant figures of speech, as found in that most underrated of ancient subjects, rhetoric. For instance, in the film *Lock, Stock and Two Smoking Barrels*, tough guy Rory Breaker threatens:

> *If you hold back anything, I'll kill ya. If you bend the truth or I think you're bending the truth, I'll kill ya. If you forget anything, I'll kill ya. In fact, you're gonna have to work very*

[15] *The Illusory Truth Effect: Exploring Implicit and Explicit Memory Influences on Consumer Judgments*, 2006

> *hard to stay alive, Nick. Now, do you understand everything I've just said? 'Cause if you don't, I'll kill ya!*

The repetition of the phrase "I'll kill ya" at the end of every sentence is an example of an old rhetoric technique called epistrophe, defined as the repetition of the same word or words at the end of successive phrases. Had Rory only said "I'll kill ya" once instead of four times, his fearful impact would have been greatly weakened.

Epistrophe has a twin, anaphora, in which the repeated words appear at the *start* of the phrase instead of at the end. The most famous example has to be "I have a dream..." by Dr. Martin Luther King, Jr.

> *I have a dream that one day this nation will rise up and live out the true meaning of its creed: "We hold these truths to be self-evident, that all men are created equal."*
>
> *I have a dream that one day on the red hills of Georgia, the sons of former slaves and the sons of former slave owners will be able to sit down together at the table of brotherhood.*
>
> *I have a dream that one day even the state of Mississippi, a state sweltering with the heat of injustice, sweltering with the heat of oppression, will be transformed into an oasis of freedom and justice.*
>
> *I have a dream that my four little children will one day live in a nation where they will not be judged by the color of their skin but by the content of their character.*

Back in the unimportant world of copywriting, Innocent, of smoothie fame, advertise, "No added sugar. No concentrates. No funny business". The repetition of the word "no" at the start of each phrase powerfully highlights their stance on impurities.

Another example of repetition in copywriting is with Apple's advertisement for their iPad mini, in which they highlight the con-

www.dbooks.org

cept of easiness: "And because it's so easy to use, it's easy to love."

Our final rhetoric technique for tasteful repetition is epizeuxis, which is the repetition of the same word a few times in a row. The most well-known example is the real estate mantra that it's all about "location, location, location."

If you want to move beyond repetition and find a list of other persuasive rhetorical techniques, let me refer you to *The Elements of Eloquence* by Mark Forsyth.

Don't Talk Yourself out of Markets

There is no point in having sales copy that (unnecessarily) turns away promising customers...after all, why would you want to artificially reduce potential market reach?

Despite the obviousness of this idea, some companies fail to see it. For instance, when I was growing up, my local shoe shop (and amplifier of gender stereotypes) placed "girly" Converse sneakers (read: bright colours) in the ladies' section and the more "masculine" styles (read: subdued colours) in the men's section. I remember really wanting to buy a pair of white Converse (one of their "girly" styles), but my easily embarrassed teenage self felt that they were off limits, somehow a bad fit for me and my precarious social standing.[16]

This shoe shop lost their chance to sell to my teenage self because they erected an artificial divide between male and female shoe styles. By imposing unnecessary divisions, you could say they shot themselves in the foot.

I bring this up in this book because certain styles of copywriting can have this same, undesired effect. Suppose that you have written sales copy that exclusively focuses on the value of a set of speakers to *sound engineers*. The photos used in the packaging, the jargon employed within the text, and the benefits highlighted ("great monitoring", "non-overdriven bass") sound great to the

[16] I am happy to announce that as an adult I proudly own a pair of white Converse sneakers.

music producer, but may alienate the (much larger) market of lay purchasers who would otherwise love the speakers.

I am not saying that your product has to be a one size fits all. Converse probably doesn't sell too many pairs of their All Stars sneakers to commercial lawyers. All I'm saying is that copywriters ought to be careful not to construct *unnecessary* boundaries.

Regarding the speaker set above, it may be difficult to reconcile the different marketing messages that speak to the sound engineer with those that chime with the layperson audiophile. If this is the case, the best course of action may be to release the same product twice, with different branding and marketing messages. This sort of thing happens more than you'd think. Just today, I noticed that a knock-off Apple laptop charger manufacturer has separate products on Amazon for the chargers for each type of Macbook—"Macbook 2016 charger", "Macbook Pro 2014 charger", "Macbook Retina 2015 charger", etc. This is the case *even though* all these Macbook models accept the *exact same charger*.

Absurd? Not at all. The point of this proliferation of charger names is to clarify beyond doubt to the casual Apple user that this manufacturer's chargers will fit whatever Macbook model the customer owns. It assures the uncertain customer that they won't need to check the fine print or worry about having ordered the wrong one.

This is marketing done at its very finest.

Keep the Scent

Suppose that one day, while minding your own business on the internet, you encounter a bicycle advert that is set against a backdrop of earthy browns and clay reds. The ad's text is written in a playful tone and brims with attitude. Underneath the main sales copy, the advert prominently promises free delivery. Intrigued, you click the ad, only to arrive on a drab and serious-looking bike-retailer website. Its colour scheme is dominated by navy blues and whites, its copy is corporate and soulless, and it makes no

www.dbooks.org

mention of free delivery anywhere. At this point you feel disoriented, unsure as to whether you clicked through to the right website at all. You hit the back button and try clicking the advert again, only to return to this place of confusion.

There's a term for this special kind of advertising problem—"losing the scent"—and it's known to hurt conversion rates for the reasons given above.

Losing the scent can also happen when the content on a web page isn't easy to connect back to the advert that brought traffic there. Website owners cannot assume that leads will have the patience to connect the dots from the information given to them. For instance, if an advert has the headline "Start a Business" but the landing page has the headline "XYZ Accelerator", then this mismatch may disorientate visitors who clicked the advert expecting to find a vanilla service for starting a normal business. These visitors mightn't bother reading farther down the page, so they will never get to the part of the page where the accelerator explains that one of their services is helping young entrepreneurs start businesses. The accelerator would probably have been better off if their landing page headline (connected to that advert) read "Start a Business". The fact that XYZ Accelerator was performing this service is only of secondary importance to the visitor. Accordingly, it should have instead been mentioned in less prominent text farther down the page.

So far we've only talked about keeping the scent between *online advertisements* and *landing pages*, but it's worth mentioning that this same need for correspondence applies at every step of the advertiser's sales funnel—from paper brochures to social media profiles to marketing emails.

Arouse Curiosity

Curiosity—the desire to know something—is a powerful motivator. Marketing techniques that leverage curiosity—say by hinting at some fact without fully revealing it—are used with good effect to turn leads into converts. To give an example of how this can

happen, look at the experiences of dating website *FindSomeone*.[17] From their homepage, it's possible for non-members to search for a nearby partner. Originally, their website displayed short profile stubs for *all* potential matches—no matter how many—and they only asked the non-member to register if they wished to read the profile in full or to contact their match.

Later on, they redesigned their website such that only the first five or so matches were shown, cutting off the rest with a big obstructive overlay that read "Discover Others Who Are Looking". This overlay was cleverly placed such that it cut through half of the fifth profile stub—you'd see this match's forehead and eyes but not their mouth. These changes, which *reduced* how much of their content this website showed to non-members, managed to increase conversion rates by 21%.[18] The explanation: Visitors were curious to see what was behind the membership wall.



Figure 11.1 – *Source:FindSomeone.co.nz*

Quora's strategy for encouraging registrations is probably based on the same idea. Visitors arriving from external websites (e.g.,

[17]https://findsomeone.co.nz/
[18]http://www.goodui.org/evidence/test021

www.dbooks.org

Google Search or Reddit) can view one page of Quora when logged out. But if this visitor clicks on any link within Quora, they'll be blocked by a registration wall. The exact design of this registration wall is telling: The question, placed at the top of the page, is visible, along with the first answer, albeit in greyed out and (often) cropped form. The reader, who will at this point become curious, will feel compelled to register to see the rest of the answer.

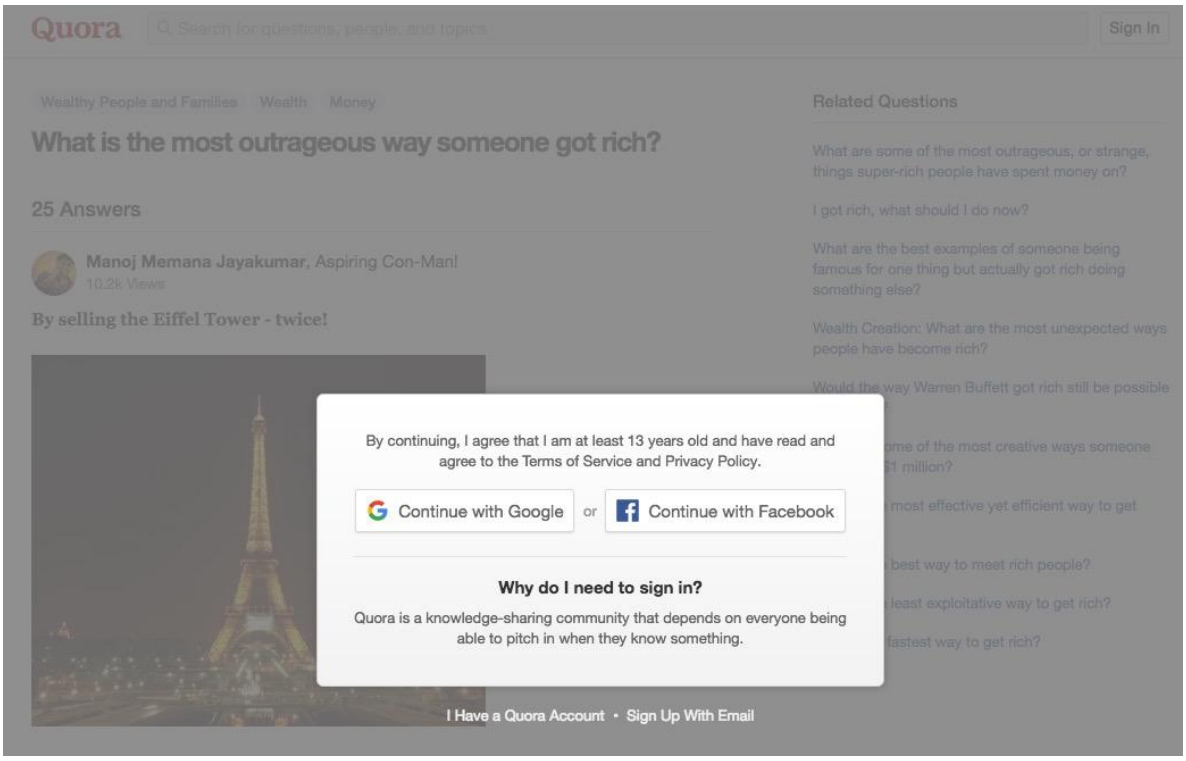


Figure 11.2 – *Source:Quora.com*

The use of curiosity is taken to an extreme by click-bait websites that intrigue with their headlines, often only to disappoint in the actual article. A traditional newspaper might report a story under the headline "Obama says gay marriage should be legal", whereas a click-bait website might report the same with "Now THIS is why I voted for Barack Obama". The traditional newspaper's headline might be descriptive and contain the right keywords for SEO, but it isn't intriguing enough to arouse the curiosity needed to get clicks on social media platforms like Facebook.

The impact of employing curiosity in the titles of articles

or videos should not be underestimated. Peter Koechley, co-founder of viral media website Upworthy, writes that one video got a 1500% increase in views when it changed its title from "Zach Wahls Speaks About Family" to "Two Lesbians Raised A Baby And This Is What They Got".[19]

So far, we've mostly talked about employing curiosity in persuading people to visit your website or register. But curiosity also plays a role *within* the one page, usually by keeping someone engaged enough to keep reading a long article or sales page. If you arouse a reader's curiosity in your headlines and opening words, then this reader will keep reading until their curiosity is satisfied. This gives you extra time to convince them to trust you or to buy from you. In a similar vein, authors of crime novels often foreshadow a dramatic scene in their opening pages, so as to hook in anyone browsing the book in a store.

Play on Negativity

Headlines that focus on negativity ("Never Do This", "Worst Problems with X") have been shown by content discovery platform Outbrain to receive a 30% higher click-through rate than those that are neutral. When negativity is compared to positivity, the effect is even more pronounced: Negative headlines had an approximately 83% higher click-through rate than positive ones ("Best Way to X" etc).[20] The cause of this difference was conjectured by Outbrain to be that "sources of negative information may be more likely to be perceived as impartial and authentic". This trust in authentic information, then, would encourage more people to click through.

Another possible explanation for this "negativity bonus" could be that human beings are psychologically primed to prefer avoiding loss to acquiring gains.[21] On this basis, sales copy that focuses on

[19] http://de.slideshare.net/fmsignal/this-title-matters-more-than-my-talk-speaker-peter-koechley-upworthy

[20] http://www.outbrain.com/blog/2013/07/headlines-when-the-best-brings-the-worst-and-the-worst-brings-the-best.html

[21] https://www.wikiwand.com/en/Loss_aversion

www.dbooks.org

the dangers of not using a product will be more compelling than copy that focuses on what the customer gains by owning it.

Fear of missing out is yet another way marketers employ negativity. ConversionXL studied how varying the decline text in a pop-up affected sign-ups for a mailing list. They found that the decline text "No Thanks, I prefer to pay full price for my clothing" generated 34% more email sign-ups than text that said "No Thanks!"[22]. This shows not only that negativity works, but also that marketers should pay as much attention to their declines-to-action as to their calls-to-action.

Anchor Numbers

Anchoring is a psychological technique that flirts strongly with the darker side of the advertising arts. The basic idea is that the mere mention of a number in an earlier, unrelated context can influence someone's judgement about what an appropriate price or order size might be. There is compelling evidence of this effect, both in the scientific community and in the business community.[23]

A good example of anchoring in the business world comes from an American restaurant, the Big Texan Steak Ranch, that mentioned a free 72-ounce steak in their marketing. (The catch for customers was that if they *didn't* finish the steak, they would have to pay for it.) Even though most customers would never take on this challenge, the figure of 72, emphasised in their marketing materials and in their menus, subtly raised customer's estimates about how much food they could handle and how much money is appropriate to spend on their dinner.

Back in the web business, anchoring can be used to make your prices seem more attractive. This could be done by mentioning the 10,000 MB of storage supplied to each account of your cloud storage platform, highlighting the 70,000 happy customers you

[22]http://conversionxl.com/popup-defense/?hvid=2EcGFw
[23]The book *Priceless* by William Poundstone has countless examples of anchoring.

have, emphasising how many thousands of pounds you had to invest in developing the product, or talking about how much your product is "worth" (as opposed to how much you sell it for, which may be much less). Placement of the anchors matters too. Often it's best to put the anchor just above where you reveal the price of your product.

Give Detailed Information

It is unlikely that a marketer would be able to sell a $9,000 guitar with two bullet points about its features. Before a customer makes a big-ticket purchase, they will want to read detailed descriptive text, so be sure to give them plenty of information about your product's benefits, features, quirky facts, origin stories, use-cases, celebrity connections, pop-culture appearances, reviews, etc. As well as all this, take the time to provide as many photos of the product as possible, from every angle and in every type of use.

Depict Yourself as Someone Similar to Your Readers

Cialdini, in his ground breaking book *Influence*, wrote that people are more likely to buy from someone who is similar to them. This similarity could apply in one of many dimensions: gender, profession, socio-economic background, race, interests, hobbies, or dreams and aspirations. The advertiser can communicate their likeness to the customer not just in their choice of photos/copy in sales pages, but also through the information they volunteer about themselves on the company's about page.

But what is an advertiser to do if, realistically speaking, they don't have a whole lot in common with their visitors? Suppose that you are an Indian teenager and that your customers are middle-aged American parrot owners. How can you come across as sufficiently similar to them for these American customers to want to buy from you (especially given that many people, though not actively racist, may be subconsciously so[24])? The strategy chosen

[24] http://www.theguardian.com/news/oliver-burkeman-s-blog/2013/aug/15/racist-sexist-you-are-more

www.dbooks.org

by the owner of ebook *Parrot Secrets* was to adopt the pen name "Nathalie—A Parrot Lover For The Last 12 Years", a persona created especially for his customers to identify with.[25] Without this in place, it is doubtful that an Indian teenager would have been able to run such a business.

[25] http://www.cringely.com/2009/03/14/parrot-secrets/

12 | Google Analytics

Why Bother with Analytics?

Guides to website ownership seemingly all recommend that you install Google Analytics. Why the universal recommendation? What will you learn from tracking how your visitors use your website, and how will this information benefit you?

- Determining the success of all your different traffic acquisition channels. Which channel brought you more traffic—social media vs email vs paid advertising vs organic vs partner websites? And how well did traffic from each of these sources convert compared to the others?
- Determining how various marketing campaigns performed within each channel. You'll be able to see that your *Halloween Treats* advertising campaign generated more sales than your *Autumn Sale* campaign.
- Providing a problem detection system for uncovering issues with your business. Assuming that you have familiarised yourself with the normal range of values for a few highly informative metrics (such as "unique visits per week" or "revenue per week"[1]), then, just by glancing at these figures every few days, you can assure yourself that your business is healthy. But if you notice that one of these figures is way off, that most likely means that you've got a serious problem somewhere which you ought to investigate…
- Learning about how your business goals (e.g., revenue, leads, sign-ups, etc.) are correlated with other attributes (e.g., the visitor's country, the product categories they browsed, whether they are visiting from a mobile or desktop device, which website they arrived from). This information helps you focus future efforts.

[1]This metric is available only with the Analytics Ecommerce plugin.

www.dbooks.org

- Understanding your website's seasonality. Through graphical analysis of past traffic and revenue, you will be able to discern seasonal patterns and be better able to project and anticipate future seasonal variation.
- Determining the most successful landing pages. When you know which ones were effective, you can study them with the goal of creating more of these kinds of successes in the future.
- Measuring which browsers and devices were most commonly used by your visitors. This enables you to decide the relative priorities to assign each platform in your graphic design and your software testing. Optimising for desktop is wasteful if 98% of your traffic comes from mobile.
- Improving the effectiveness of Google AdWords advertising. With a bit of configuration, you can see how much profit you are earning (on average) for each click on a given paid keyword. This, in turn, will inform your bidding in Google AdWords, ensuring that you are always profitable in your advertising.
- Gathering demographic data, such as what countries and cities your customers are based in. This is especially relevant when geography influences what products you stock or what languages you should present your website in.
- Measuring the snappiness of your website. Identifying slow pages can give you some "quick" gains in SEO and conversion optimisation.
- Discovering which pages of your website are most engaging, as evidenced by page view counts, repeat visitor counts, or average times on page. This knowledge could inform website redesigns, for example by prompting you to link to these beloved pages from the home page.
- Alerting you of web pages that are accidentally publicly available. For example, if some part of your website was supposed to be behind an admin-only password wall, then a Google Analytics report that says hundreds of people saw that content indicates that you've got a security bug.
- Knowing what Google Search queries your website appears for, along with your website's rank in the results page and

the click-through rate your entry achieves.[2]
- Learning what your visitors search for within your own website (via your on-site search box). This could prompt you to start stocking products which you don't yet sell, but which your customers are already looking for.
- Pinpointing the stage within your funnels where the most leakage occurs, and thereby empowering you to focus conversion optimisation efforts on these problem points. For example, do a higher percentage of people drop out at the payment stage or the registration stage?
- Identifying which pages have high exit rates or high bounce rates, either of which would indicate that a page is scaring away traffic and needs some work.
- Stroking of ego. Seeing tens of daily websites hits grow to hundreds and then to thousands is gratifying—and this indulgence is a healthy one, thanks to its motivational effect.

Workflow Tips

The sanity-preserving Analytics workflow: Form plain English sentences backed by hard data

Google Analytics is daunting: There are so many report screens with so many variables and tiny little numbers that it's all too easy to drown in data. But in spite of this overwhelming feeling, it is nevertheless still worth wading in and exploring. Indeed, some of the best insights can come from wandering around the data without a plan.

There is a nice trick to help you regain a feeling of control during these exploratory sessions: Approach each session with the express goal of forming a few sentences about your business, as if you were going to present some curious facts and observations to your friends over lunch. This means you want to end up with sentences like "Fifteen percent of our web visitors come from Germany, yet none of them end up buying our products" or "Our pricing web page has double the amount of clicks as any other

[2] This assumes you have Google Search Analytics installed in a connected Google Search Console account.

www.dbooks.org

page we display in our global navigation" or "MS Windows phone users are four times as likely as users of other platforms to leave the website within five seconds of arriving". When data is condensed into English like this, it becomes more memorable and far more actionable.

Navigate Analytics by searching

It's always a pain in the ass to find a specific report in Google Analytics. Often the one you need lies nested—or doubled nested—somewhere deep in the left-hand navigation panel.

Instead of straining to find a report by trial-and-error browsing, it's much easier to search for it in the upper-left-side search bar titled "Find Reports & More".

This search works like magic, and it will save you a whole lot of time.

Connect to Other Google Products

AdWords

Background: I believe that anyone running an online business should, at the very least, try to find themselves a profitable stream of traffic through Google AdWords. I deal with the specifics of paid advertising elsewhere in this book, but I want to mention the importance of associating your AdWords and Analytics accounts here.

Why do this? One of the biggest reasons is that you'll get access to a Google Analytics column called ROAS ("Return On Advertiser Spend"), which is basically a god figure for advertisers. ROAS is calculated by dividing the revenue generated by your Google AdWords traffic[3] by the cost you incurred in buying that traffic.[4] The reason this figure is so great is because it sums up the profitability of your advertising by AdWords campaign, by keyword, or even by landing page.

[3]This assumes you have e-commerce tracking for Analytics configured.
[4]This data is imported from Google AdWords.

An ROAS of 100% means that your revenue was equal to your advertising spend. You'll want an ROAS of at least 100%, if not significantly more, so as to account for your other expenses in addition to advertising costs. If you know your average profit margins, then you'll be able to calculate a minimally acceptable ROAS figure, and then ensure that you only advertise whenever this figure is exceeded.

You can access the ROAS column by viewing any of the Analytics reports filed under Acquisition > AdWords and then clicking on the "Clicks" tab within the "Explorer" bar. After doing this, the ROAS column (and a few other related ones) will appear in the main data table.

Yet another good reason to associate your AdWords and Analytics accounts is to unify conversion tracking between the two platforms. For readers not so familiar with these platforms, you should know that both have their own way of tracking conversions: AdWords has its "conversions" feature and Analytics has its "goals"/"e-commerce" tracking features. Theoretically, you could use the conversion tracking provided on both platforms, but this isn't recommended. For one, the AdWords conversion tracking features are not as flexible as the Analytics ones, and this will make integration more difficult for your tech team. Moreover, simultaneously tracking conversions in AdWords and Analytics creates the risk that the figures recorded in one platform start diverging from those in the other. This risk is especially heightened after you make changes to your website that require you to update your conversion tracking (it's possible to forget to update things for one of the two platforms). Unnoticed problems like these could lead to you getting different conversion numbers depending on where you look. Save yourself the hassle and skip the AdWords conversion features completely—it's easier to set up Analytics goals and link them to your AdWords account.

Convinced? Great! A guide to associating these two accounts can be found by following the link in the footnotes.[5]

[5]https://support.google.com/adwords/answer/1704341?hl=en

www.dbooks.org

Search Console

Google Search Console (formerly known as Webmaster Tools) helps you monitor your site's presence within Google Search results, helping you answer questions like "which queries caused my site to appear in search results?" or "how highly do I rank in the results for this query?" or "how many people *clicked through* to my website after seeing it mentioned within the search results?"

Although this information is already available within the Search Console[6] platform, it is nevertheless helpful to have it available in Google Analytics so you have everything in one place. This is done by asking Google to associate your Search Console account from your Analytics admin page.[7] Please be aware, though, that because of privacy concerns, Google Analytics places heavy limits on how you can associate your search results data with other metrics in Analytics. If you run into a block, it probably isn't a bug...

Location of Snippet in HTML

Your Google Analytics code should appear inside the <head> section of each page's HTML. This is because code which appears higher up in HTML gets executed sooner, and code that gets executed sooner is less likely to be skipped if a visitor happens to navigate to another page before the current one has finished loading. By having the Analytics code high up in the page, you decrease the chances that Analytics will underreport hits or conversions.

Implementing this isn't as simple as copying and pasting the Analytics snippet high into your website's <head> and then calling it a day. Problems appear when you have an advanced Analytics installation that requires you to send additional data to Google (e.g., ecommerce revenue data, events, custom dimensions, or User IDs). In order for this extra data to get sent, it must be already prepared and available within the website <head>, and this

[6]Under Search Traffic > Search Analytics
[7]Details here: https://support.google.com/webmasters/answer/1120006?hl=en

requirement can have major consequences for your website's architecture or caching layer.

Cross-Browser and Cross-Device Tracking (User ID)

Gone are the days when people interact with your business from their one and only computer. Instead, they first encounter your website on their phone, next download your app, and then finally complete a purchase on their laptop a week or two later.

By default, Google Analytics records these three interactions as coming from three distinct users. But this is a distortion of what actually happens. Without any corrective intervention, Google Analytics will give the impression that the traffic from your mobile website and your app was totally worthless, whereas the traffic from your laptop users was gold. But the truth is that the mobile traffic was crucial for making first contact, and there would be no eventual purchases from laptop devices without this step.

Why does Google Analytics falter in this way? This comes down to the technicalities of how Google ties browser sessions together into overall user stories. Whenever someone visits a website that has Google Analytics installed, Google drops a cookie on that visitor's browser. Should that person happen to visit the website again, Google will notice that a cookie was already dropped on the browser, and they can use this data to link the earlier session to the current session. But there is a big limitation to all this: Due to how cookie technology is designed, cookies are specific to one browser (i.e., they are neither shared with other browsers running on the same device nor with browsers running on other devices). This means that Google Analytics is unable to stitch sessions together which occurred on multiple devices or on multiple browsers. And to further complicate the picture, even the cookies within the one browser can expire or get deleted.[8]

Google Analytics has a feature called User ID that can overcome

[8]By default, Google Analytics cookies only expire after two years. Please consult this reference: https://developers.google.com/analytics/devguides/collection/analyticsjs/cookie-usage?csw=1#analyticsjs

www.dbooks.org

these difficulties and help you stitch together sessions occurring on multiple browsers/devices that would otherwise appear separate. To use this feature, your website must generate a unique ID for each user and send that ID to Google Analytics. If tracked traffic from your iPhone app sends the same User ID as the traffic from your regular website, Google Analytics can see that both incidents belong to the one story. Be careful that you don't send Google a User ID that contains information that would enable them to personally identify this person (e.g., a username or email address). Instead, send an identifier only known to you (e.g., an ID number you generate in your database).

There is a proviso for this feature to work: You have to be able to reliably create, store, and assign a unique ID to your users *during every session*. The most typical implementation of this is to tie in the Analytics User ID to whatever ID you use internally to manage logged-in user accounts.[9] This makes sense because it leverages existing login and database functionality, but it has the downside that your tracking will not account for visitors who don't log in during a particular session. (However, this setup *does* manage to capture visitors who begin a session while logged out but log in later during the same session.[10])

After you set up User ID tracking on Google Analytics,[11] you will be encouraged to create a separate "view" of your Analytics data that will be limited exclusively to the visitors for whom you have User IDs.[12] When you switch your Analytics reports to this view, you'll see all your usual reports, along with a few new ones which would be otherwise unavailable. These are filed under Audiences

[9] An implementation that doesn't depend on log-in functionality could be achieved, for example, by sending out bespoke URLs to each individual lead and then having the software assign unique User IDs based on which of these bespoke URLs was used to access the website.

[10] Through an Analytics feature called Session Unification, hits collected before the User ID is assigned can be later associated with the User ID once it becomes available. This helps you analyse the behaviours that *lead up to* the assigning of a User ID (e.g., actions taken before creating an account on your website). Read more: https://support.google.com/analytics/answer/4574780?hl=en

[11] Installation details: https://support.google.com/analytics/answer/3123666

[12] Frustratingly, Google Analytics disables goals and ecommerce tracking on new views, so be sure to re-enable them under Admin > Account > Property > View > Goals (/ E-commerce Settings).

> Cross Device.

- The Device Overlap report lets you answer questions like "what is the conversion rate for customers that interact with my business through the combination of devices X, Y, and Z?" (You might learn that you have a 1.2% conversion rate for users that only use desktop and 1.3% for users that only use mobile, but a 3.4% conversion rate with users that accessed your website on both desktop and mobile. This would indicate that you should advertise on both platforms simultaneously.)
- The Device Acquisition report helps you understand the value of various "first contact" devices. You can answer questions like "how much revenue do I eventually earn when a user first encounters my business using device X?" If you know that the users who find you on mobile are the most likely to eventually buy, then you should focus your "fresh leads" advertising budget on this more profitable group.
- The Device Paths report lists which devices were used (in what order) within the steps taken before (or after) some goal completion, event, or ecommerce transaction. This information could let a SaaS (Software as a Service) company know which devices to prioritise for each goal type.

A word of warning: User ID tracking isn't and cannot be retrospective, so be sure to have it set up before you launch your website.

Subdomain Blues

Conflation of traffic with the same relative paths on different subdomains

By default, Google Analytics records only the relative path of a hit (e.g., "/about_us") instead of the full path with the host name included (e.g., "www.mysite.com/about_us"). This becomes problematic when your website is hosted across multiple subdomains. As far as a Google Analytics account configured to relative paths can see, hits to "blog.mysite.com/index.html" are indistinguishable from hits to "www.mysite.com/index.html". In its eyes, the

www.dbooks.org

URL "/index.html" was simply visited twice.

This is obviously not what you want, and this sort of conflation can invalidate your reported data by incorrectly bucketing together hits to distinct entities. The solution to this conundrum is to configure Google Analytics to record the full URL instead of just the relative path name. This is done with a filter that rewrites the "page" dimension within Google Analytics so as to include the host name. [13] Be warned though: Adding this filter to an existing account will invalidate existing Google Analytics goals (or other filters) that are programmed to trigger on exact matches of URLs. For example, a goal configured to trigger when the page exactly matches "/thanks_for_buying" will no longer trigger when the page data becomes "www.mysite.com/thanks_for_buying" because the additional host name information means a character-for-character perfect match no longer exists. This means that your Analytics installation will stop recording goal completions until you update your goal definition to work in harmony with full URL names.

Session data resetting across subdomains

Whenever a visitor navigates to a different subdomain on your website, their session data will be reset, meaning that, as far as Analytics is concerned, a separate visitor arrived, and this "new" visitor was referred to your website by your other subdomain (instead of by the true outside source like blog X or social network Y or advertising campaign Z). This session resetting causes all sorts of distortions in your reported data. For example, it will decrease the reported session length of each visit, it will clobber over information about which marketing channel originally delivered the visitor to you in the first place, and it will mess up analysis of the value of each marketing channel.

The fix to theses ills comes in the form of configuring Referral Exclusion Lists for your various domains[14] and integrating Google's

[13] https://developers.google.com/analytics/devguides/collection/gajs/gaTrackingSite#profilesKey

[14] Admin > Tracking Info > Referral Exclusion List. See guide: https://support.google.com/analytics/answer/2795830?hl=en

autolink functionality.[15]

Cleanup Noise, Chaff, and Distortion

Lowercase Request URL

In the world of URLs, character case matters. By way of example, the following three URLs count as separate entities:

```
/products/iphone
/products/IPHONE
/products/Iphone
```

And because these three are considered separate entities, Google Analytics dutifully records them all as distinct items in its reports. This would be a good thing if your website showed different content depending on the case of the URL. But such a website architecture is rarely done (and is terribly confusing when it is done), so I would argue that Google Analytics' default treatment is going to spell trouble for the typical webmaster. This is because its effect is to pick up URLs entered incorrectly by your users (or linkers) and cause their tracking data to disperse across the variously cased URLs in your reports, ultimately causing your website's performance to be more difficult to analyse.[16]

The solution is to merge all these separated entries into a single Google Analytics entry by creating a custom Google Analytics filter that lowercases the request URL.[17]

That said, this problem should really be rectified within your web application's software layer. Your website should only respond with HTTP status code 200 to *lowercased* URLs; when someone enters an uppercased variant, your website's router should *redirect* (i.e., HTTP status code 301) to the lowercase version, or just not respond at all. But sometimes making this software change

[15] https://developers.google.com/analytics/devguides/collection/analyticsjs/linker

[16] Of course, this will only be a problem if your website accepts URLs in various cases (i.e., responds with HTTP status code 200 instead of 404ing or redirecting).

[17] Do this by entering Admin > Views > Filters. Within the Filter Type field, select "Custom", then select "Lowercase", and then select as the Filter Field "Request URI".

www.dbooks.org

isn't possible or practical, and in these cases the Analytics filter in the preceding paragraph will get you out of a fix.

Lowercase all UTM parameters (URL Builder parameters)

UTM parameters smuggle extra data into your Google Analytics reports by appending information to the end of your URLs. These parameters, which you'll either set yourself or have platforms like Google AdWords set automatically for you, contain information about the source of the traffic (e.g., marketing channel, advertising campaign, or CPC keyword that generated the traffic).

I bring up UTM parameters here because there is a danger that your UTM parameters may appear lowercased in one advertising campaign and upper/mixed cased elsewhere (e.g., utm_source=email in one place and utm_source=Email elsewhere). This would cause Google Analytics entries for the email source to be split across two places ("email" and "Email"), making analysis annoying to carry out and dramatically more error-prone.

While it's certainly possible to instigate company policies which demand that you and your staff always lowercase UTM parameters, realistically speaking, you can't guarantee that this policy will be consistently applied—after all, human memory is oh so fallible. A better solution would be to create filters within Google Analytics that automatically lowercase each of the UTM parameters, thereby ensuring that this problem can never occur.[18]

Remove sorting/pagination/noisy parameters

In some quarters of your website, your software might automatically append parameters to the URL. The classic example is when you have a listing page with the option to sort the listed products by price, size, or some other attribute. A common implementation for this feature involves redirecting users to a modified URL

[18]This process of creating a filter is the same as in the previous footnote, except that you'll now need five separate filters, one for each of the UTM parameters: Campaign Medium, Campaign Source, Campaign Content, Campaign Term, and Campaign Name.

which has been updated to include these sort parameters. (The reasoning for such a URL design is that appending the sort parameters to the URL enables the user to bookmark or share that particular view of the page.)

Here's an example of one such URL: A student shopping for medicine textbooks sorts by price in ascending direction, and you can see that these sorting instructions are encoded into the end of the URL:

```
/taxons/medicine?direction=asc&sort=price
```

Google Analytics (as with the rest of the internet) treats each URL as unique, so, within your reports, you will see separate entries for every combination of sorting parameters. This means that your report could contain distinct entries for

```
/taxons/medicine?direction=asc&sort=price
/taxons/medicine?direction=desc&sort=price
/taxons/medicine?direction=asc&sort=institution
/taxons/medicine?direction=asc&sort=name
```

But this is not what you want, from an analytical point of view, because it means that your data about the "/taxons/medicine" page will be dispersed across four entries and correspondingly awkward to work with. From the business point of view, each of the above four entries contains *the same* content (i.e., medicine textbooks). The various entries in the Analytics report only represent *different orderings* of this same content. It is more useful for you, the marketer, to bucket all the above into a single Analytics entry for

```
/taxons/medicine
```

Luckily for us, it's possible to configure Google Analytics to automatically collapse URLs that differ only in sort (/filter/pagination parameters) into their proper catch-all URLs. To do this, you simply tell Google Analytics the names of the URL parameters which you want it to ignore. This is done through the Exclude URL parameters feature.[19] In the above sorting example, we'd tell

[19] Go to Admin > View Settings > Exclude URL parameters. If you are enabling this feature for

www.dbooks.org

Google to ignore the "direction" and "sort" parameters, and this instruction would collapse the four entries into one for "/taxons/medicine".

Some caution is advised in setting up these filters because they apply not just to parameters in individual URLs, but also to parameters contained in *any* URL across your whole website. If you aren't careful, you might unintentionally collapse together URLs elsewhere on your website that really ought to appear distinctly in your reports. For example, what if "sort" was used to mean "type" or "variety" in some other part of your website (e.g., for product variant pages like "/products/Xlaptop?sort=15-inch-screen" and "/products/Xlaptop?sort=13-inch-screen")? You may well want these two variants to appear as distinct entities in your reports.

The best way to avoid these kinds of unintentional clobbering problems is to be consistent with the parameter names used on your website for sorting, filtering, or paginating content. This is something that should be thought about up front when programming and designing the URL surface area of the website. If the parameter name "sort" is used for sorting data, treat this as reserved and ensure it isn't used elsewhere to mean something else. (Ditto for "page" and other parameters).

Above, I talked about filtering out parameters for reordering content (such as sorting and filtering). But there may be other kinds of pesky parameters that add nothing to reporting, and that you'll consequently want to exclude. Indeed, anything that's noisy, unimportant, or that leaks sensitive information (such as customer names) ought to be expelled from the URL. This means that you will probably want to strip out parameters like the following:

---

an existing website, you might want to do a full audit of all the existing URL parameters ever recorded by your Analytics installation for this website. Do this by switching to the all-time view of your data and then going to Filter Behaviour > Site Content > All Pages. Next, search for the telltale symbol that syntactically indicates the start of a parameter listing in URLs: "?". The results of your search will show all the URLs stored in your Analytics view that contain query string parameters (e.g., "?sort="), and you can browse through these and locate the ones you'd like to filter out from now on.

```
order\_token=
utf8=
affiliate\_token=
customer\_email=
```

Exclude bots

The internet is inhabited not just by humans, but also by bots. These robotic ghosts will visit your website day and night and thereby inflate your view counts. While this is great for bragging rights and for impressing hopelessly naïve investors, ultimately, the view count boost is a Pyrrhic victory because these bots will taint and distort all your other statistics, like traffic acquisition channels, conversion rates, bounce rates, and session durations. Your reports, now muddied, will leave you with incorrect impressions about your website, and any intelligence you once gleaned from studying your metrics will now be lost—unless you configured Google Analytics to exclude bots.[20]

Exclude internal IP address from your live tracker

Whenever you visit your own live website, you skew your Google Analytics results ever so slightly. This happens because your activities do not represent those of the typical user. The average user, after all, does not spend full days proofreading, tinkering with the CSS, or checking each page's performance in Chrome DevTools.

More problematic are the effects of internal automated tools that interact with your live website. I have in mind here cache-warmers or stress-testing suites, both of which deliver a lot of unrepresentative hits.

For these and other reasons, you should exclude internal traffic from Google Analytics reports. Assuming you work from a static IP address, you can sort this out by setting up some IP filters.[21]

[20] Exclude bots in Admin > View Settings.
[21] Google have a guide on doing this: https://support.google.com/analytics/answer/1034840?hl=en

www.dbooks.org

Set up a separate web property in Analytics for staging/testing

Absent any special precautions, all your activity on testing/local/development versions of your website will pollute your live Google Analytics data. When your programmer adds something to the cart 20 times during a bug fix, your business statistics on Google Analytics will be similarly inflated.

There are a few ways around this, but the most useful is to send data from test versions of the website to a separate Google Analytics tracking profile.[22] This has two advantages. Firstly, it stops you polluting your live data. Secondly, the new "test" profile will help you verify—before deploying—that Google Analytics will continue working after you deploy your most recent changes, helping you to catch potential reporting bugs before they do irreversible damage.

Exclude unwanted referrers

Every time a visitor leaves your website only to later return, Analytics tracks a new session, associating it with a new "referrer" website. Normally this is perfectly sensible and correct, but there are a few isolated circumstances when recording a return visit in this default manner will distort your reporting data. Suppose you are using an off-site payment provider, like Paypal. Someone comes to your website via a Facebook advertisement, adds a product to the cart, checks out, gets redirected to pay on Paypal, then returns to your website and sees the thank-you page. By default, Google Analytics views this last hit as a brand new session, referred by Paypal.com. If your website displays a conversion pixel on the thank-you page (as many websites do), then this conversion will be attributed to traffic *referred by Paypal* instead of traffic referred by your Facebook advertisements. This makes it look like Paypal.com is a major referrer of converting traffic, when in reality it was Facebook that brought home the bacon. The worst part of it all is that the traffic coming from Facebook advertising is reported as accumulating in nothing but cart

[22] Do this by creating a separate web property in Google Analytics that has its own unique profile ID. Next, configure your website software to send tracking data occurring in the test version of your website to this distinct profile ID.

abandons, occurring on the payment page.

The best solution to this distortion is not to use off-site redirects within your flows. But if it isn't possible to remove these, you can still prevent this reporting problem by adding a Referral Exclusion for paypal.com.[23] Setting this exclusion causes the user session to continue after the visitor returns from Paypal, meaning that Google Analytics will attribute the conversion to the original referring source, Facebook, instead of to Paypal.

Analytics and Social Media

Social interactions

Google Analytics has a feature to record social network interactions such as Facebook likes/shares or Twitter tweets. Why are these social media interactions worth tracking? Because Google Analytics gives you convenient reports about your social media performance. This convenience stems from you being able to see—in the one Analytics report—how each of your pages fared. This beats flicking from page to page, especially if there are tens of thousands of pages on your website. Furthermore, manually checking the counters on each page is increasingly becoming impossible because social media platforms such as Twitter are starting to remove the counters from their on-site widgets. When these restrictions are in place, you'll have no idea about how often some piece of your content was shared–unless you are tracking it with Analytics.

The main advantage of tracking shares is knowing what type of content is striking a chord with your readership so you can create more content like it in the future. There's also a smaller advantage from the point of view of website design: Knowing which network share buttons get used and which ones don't will prompt you to remove the underperforming buttons so as to reduce website clutter.

The Analytics report isn't totally comprehensive though: It only

[23] Admin > Tracking Info > Referral Exclusion List

www.dbooks.org

tracks interactions that occur when someone uses the social media plugin installed on your website. The report knows nothing about social media interactions that happen elsewhere (e.g., retweets that originate on Twitter or manual copy-and-pastes of your URL into a Facebook status update).

Social interaction tracking is implemented using Javascript that subscribes to social plugin events and forwards this information to Google Analytics.[24] For marketers who don't want to mess with setting this Javascript up manually, there is a lazy option (the Autotrack plugin) that requires no configuration and is available to anyone using the official Twitter and Facebook buttons.[25]

Events for Remarketing

The fact that someone is sharing your website on social media shows that they are highly engaged with your brand and appreciate your work. What's more, they are part of the "sharer-erati", the active and vocal subset of internet users predisposed to sharing content. Who better than them to remarket future content to?

The logical thing would be to build a remarketing audience in order to target those that Google Analytics records as having shared through their Social Interactions feature (as we set up above). Unfortunately, it isn't possible to create an audience from this data, so we must instead create a Google Analytics event whenever someone shares content, and then build our lists up out of these lists.

Custom Dimensions

Even without any customisation, Google Analytics comes with a great many built-in dimensions for analysing and segmenting your data (e.g., country of visitor, referring website). But despite all this, your efforts at gaining insight might nevertheless be hampered by the lack of some crucial piece of information specific to your business. Custom dimensions fill this gap by giving you the

[24] https://developers.google.com/analytics/devguides/collection/analyticsjs/social-interactions#overview

[25] https://github.com/googleanalytics/autotrack

option to send Google Analytics additional information for your reports.

Let's look at some use cases that demonstrate how you might use this feature. Imagine you own a newspaper website and have a bunch of different columnists writing articles for your paper. Because talent is never evenly distributed, some authors will be better at engaging readers than others, so it's important for you to know who these star performers are. By default, Google Analytics only records the URL names, but this doesn't do you any good because your website makes no mention of the author's name in the URL (you want to keep your URLs short and sweet instead of turning them into overladen mules for hauling information to Google Analytics[26]). Essentially there's no practical way for you to know how engaging the public finds each of your authors—that is, unless you add a custom dimension for "author name" that gets sent to Google Analytics along with every tracked hit. With this in place, you can start segmenting and studying your data along these lines.

Here's a second example: Suppose you own an ecommerce store that sells books in various genres ("humour", "politics", "business"). Without any custom dimension for "book genre", you would be unable to tell how each genre performs relative to the others in your Analytics reports.

Illuminating Analytics Blind Spots

Virtual pageviews

Online marketers often have the impression that including the Google Analytics snippet into their website templates is enough to track each and every page view occurring across their websites. Unfortunately, it's not that simple.

Let's take a demonstrative example: Suppose you have a "contact us" form found at the URL "/contact_us" (via a HTTP GET

[26] Not to mention that storing information in the URL doesn't scale…what if there were *10* dimensions you wished to send to Analytics?

www.dbooks.org

request). Your website is designed such that whenever a question is submitted through this form (via a HTTP POST request), the customer gets redirected back to the "/contact_us" page, where they see the same "contact us" form again, except this time there is a little flash message at the top of the screen that reads "Your question was successfully received". From a human and common sense point of view, it's clear that two distinct things happened: 1) Someone viewed the "contact us" form and 2) someone sent a question through that form. But from the point of view of Google Analytics, which only has the URL information to go by, it just sees that the "/contact_us" URL was visited twice; it cannot know that the second hit, recorded after the question was sent, represented a qualitatively different event. Information is thus lost.

This sort of inaccurate reporting doesn't just happen with respect to forms that submit data to the same URL as where the form was found—it also affects heavily AJAXed components, like one-page checkout processes. Not only that, but Google Analytics usually can't track hits to non-HTML content on your website, such as MP3s or downloadable PDFs. This is because these sorts of content are usually served outside the frame of your website template, and so your Google Analytics tracker is unavailable to register hits.

The most generally applicable solution to these problems comes in the form of virtual pageviews,[27] a Google Analytics feature whereby you insert Javascript code that registers page views for pages that "don't exist" (from the URL-centric point of view)—thus explaining the "*virtual*" in this feature's name. Returning to our original example, after someone submits a question on our contact form, we could have our code send the flash message via Javascript (so as not to create a duplicate hit to "/contact_us"), and then we could record a virtual pageview for "/contact_us/submitted" (even though the URL "/contact_us/submitted" doesn't actually exist on our website).

[27] https://developers.google.com/analytics/devguides/collection/analyticsjs/single-page-applications

Those with single-page applications will, no doubt, be daunted by the idea of writing exhaustive code for tracking virtual pageviews. Luckily, Google maintain a Javascript library called Autotrack that might be able to take care of this job for you.[28] (Alternatively, some frameworks for single-page applications take care of this problem for you).

Track important occurrences that don't correspond with page views

Imagine a subscription business (like a web hosting company) that bills their customers monthly for access to their services. During the lead up to a subscription, the customer-to-be would have visited the company's website to provide their credit card details and initiate the initial billing, with the consequence that this first conversion can be tracked as usual over in Google Analytics. But during the subsequent months there is no further need for this customer to revisit the website; they can simply enjoy their web hosting service and get on with their life.

This scenario presents a tracking challenge. Google Analytics normally only sees occurrences that correspond with explicit user activity on a website or app. As a result of this blinkered vision, any activity that happens behind the scenes will not be recorded. This leaves certain metrics, such as the accumulated revenue generated through monthly billing cycles, invisible to Analytics reports. Without these data points, Analytics reports risk presenting an incomplete or even misleading picture. For instance, patchy data might suggest that advertising campaign A ("half-price basic hosting for the first three months") radically outperformed campaign B ("superior reliability and data privacy") in terms of generating revenue. But maybe it turns out that the customers won through the first campaign are fickle and leave as soon as the discount expires, whereas those acquired in the second campaign are more liable to stick around for a longer time since they bought on the basis of reliability and privacy. Because the web hosting business is all about *lifetime* revenue per customer, their goal is to optimise for long-term contracts. But with-

[28] https://github.com/googleanalytics/autotrack (Use of this plug-in assumes that your application uses the History API to update URLs.)

www.dbooks.org

out any information about subsequent billing cycles (or churn) available in their Analytics reports, this analysis task would be near impossible.

There is a solution to this conundrum: The Analytics Measurement Protocol. Roughly speaking, this Analytics feature gives your programmers the power to send tracking data about any one of your customers to Google Analytics *at any time whatsoever*—even if there has been no corresponding bout of web activity. Tracking information could be sent, for example, as part of an automatic billing process that happens once a month behind the scenes in your backend server.

You may be curious as to how Google Analytics can associate the data sent by your backend server with the regular tracking activity that occurs on the frontend. This is made possible whenever the website is programmed such that both the backend and frontend tracking systems pass Google the same User ID information, which Google then uses to stitch the two streams of events together. (This User ID is also the mechanism used to stitch together sessions occurring on different devices and browsers.)

Tracking with the Analytics Measurement Protocol has use cases beyond those of the company selling subscriptions, as we described above. Consider, for example, a lawyer specialising in workplace injuries who decides to start advertising online. Far away on a distant website, some lead sees this lawyer's adverts and clicks upon one of them. If the lead is impressed by what he or she reads on the lawyer's website, then the lead will send the lawyer a question through the website "contact us" form. The lawyer reads this message and phones the lead back. If this first conversation goes well, the lead eventually engages the lawyer's services.

Sales processes like the above cannot be comprehensively tracked with a standard installation of Google Analytics. As far as Analytics is concerned, any activity occurring after the "contact us" form never happened; these events are completely invisible, completely under the radar. This abrupt cut-off in tracking

leaves the lawyer's data open to similar kinds of misinterpretations, as we saw above with respect to the subscription business. For example, our lawyer, seeing that a certain low-cost advertising campaign caused many leads to send her messages, might end up diverting all her budget to these adverts in order to receive even more messages for her money in future. But what if, unbeknownst to her, the leads she gets from these advertising campaigns rarely buy legal services and instead only want to waste her time by asking for free advice? As it stands, this question cannot be answered; but had our lawyer integrated elements of the Analytics Measurement Protocol into her workflow, then she could simply read the answer off a report. (Exact implementation details for this might involve building an admin-only button that sends over tracking info to Google whenever some lead or other signs a contract for services.)

Collect Your On-Site Search Queries

Searches through your on-site search engine reveal useful information about what your visitors are looking for. The enterprising website owner will harvest these clues and use them to inform future inventory purchasing choices, future website design (e.g., moving commonly searched-for items to the home page), or future online advertising campaigns (e.g., mentioning the most commonly searched-for items within the advertisement copy).

With Google Analytics, you have no need to build custom server-side machinery to harvest this data; instead, you can configure Google Analytics to record all on-site searches for you.[29] Once configured, you can see what searches were most often made by your visitors, then sort and filter with respect to any other data point recorded in Analytics.[30]

One proviso: For this feature to work, your website architec-

[29] Install on-site search monitoring by visiting Admin > View Settings > and then, within the box "Query parameters", entering the name of the parameter that your website uses to store the keyword of the search query. For example, in the following URL, the parameter would be "keyword": /search?**keyword**=philosophy

[30] View the report by visiting Reporting > Behavior > Site Search > Search Terms.

www.dbooks.org

ture must display its search results on a URL which contains the search query parameter.

/search_results?keyword=philosophy

This is instead of an alternative implementation which might swallow the query parameter silently and display results on a generic URL like:

/search_results

Our initial architectural suggestion, which included the search query within the search results URL, makes sense for usability reasons too. Specifically, it allows visitors to link to particular search results pages, to bookmark them, and to send them to their friends or customer support.

Measuring Success

Setup goals

Google Analytics has a feature called Goals which you can configure to correspond to business targets like signups, purchases, leads, or any other measurable you want to improve. With Goals set up, you can study how pretty much any other Google Analytics dimension influences goal completions. For example, you will be able to tell how different marketing channels, visitor countries, or months of the year affect your conversion rates, enabling you to adapt accordingly.

There are a few different ways to configure Analytics to count a goal as having been completed, the simplest being noting when the user has reached some particular URL. To avail yourself of this simplicity, it helps to have designed your website such that goal completions correspond to URLs that are shown one time (and one time only) during some particular user journey. For example, this might entail redirecting a customer to a “/thanks_for_ordering” page immediately after payment has been made.

Track ecommerce transactions

Probably the most important business metric you'll want to keep tabs on is ecommerce transactions (i.e., orders). But instead of creating Analytics Goals to track these transactions, you are better advised to implement Google's special ecommerce tracking plugin.[31] Essentially, this plugin lets you record detailed info about what products were bought, in what numbers, and how much revenue their sale generated. As with goals, you can associate this data with pretty much anything else in Google Analytics, enabling you to correlate revenue with, say, landing page, marketing channel, or device used to browse website. This is analytical optimisation at its best.

Tag URLs

Imagine an ebook author who, during her launch week, sends out two different emails to her mailing list of 5,000 fans asking them to buy her book through a link in the email. At the end of the week, the author sees that she sold 1,000 books in total. However, she also learns that 500 people on her mailing list unsubscribed. The author is left with a few burning questions: Which of her two emails was better at triggering sales, and which one led to more unsubscribes? With a default Analytics setup, she cannot answer these questions since the traffic from both emails will appear lumped together under the one bucket: "direct" traffic. Not only will it be impossible for the author to distinguish the traffic coming from each of her emails, but it will also be impossible for her to distinguish email traffic from traffic arriving when a user types her website's URL directly into their browser. What a mess! Luckily, there is a way out of this conundrum: Our author could have *tagged* the URLs she sent out in her emails, and then Analytics could pick up on these tags and partition the traffic accordingly.

Before we explain how exactly URL tagging works, I'd like to quickly explain another important use case for this technology: that of tracking the *root* referrer of traffic, as opposed to the *pre-*

[31] https://developers.google.com/analytics/devguides/collection/analyticsjs/ecommerce

www.dbooks.org

vious referrer. An example makes this idea clearer: Imagine our book author posts to her Facebook page a witty post containing a link to her website. This post ends up getting shared hundreds of times on Facebook, and some of these shares end up getting tweeted. Now imagine a random Twitter user sees a tweeted version of the original Facebook post and clicks on the link within, bringing them to her book's website. The *previous* referrer (i.e., the directly immediate one) is twitter.com, and that's all the information that would be available within a default Analytics report. But with URL tagging, the author would be able to attribute this visitor's arrival as being *ultimately* attributed to that earlier Facebook post. In other words, URL tagging can be used to preserve information about the *root* referrer.

What is URL tagging? In abstract, it's an information-preserving hack that stows away valuable data within the URL. This info can later be picked up and unpacked by Analytics. Instead of sending out generic links to some page on your website (e.g., to https://www.example.com/buy_book), you leave slightly different versions of the link in different places (e.g., one version of the link is given in your first mailing list email, another version is given in the second, and a third version is given in your Facebook social media campaign). These links all lead to same page on your website (i.e., to the book buying page), but they differ in that each is dressed up with different extra titbits of information called "tags", which happen to be stored via query parameters appended to the end of the URL. Here's an example with just one parameter appended: https://www.example.com/buy_book?utm_campaign=newsletter-jan-8.

Most typically, the tags used are the five popularised through Google's URL Builder (we'll describe these below).[32]

Someone who has followed a tagged URL to your website will ask your server not for the standard URL (e.g. https://www.example.com/buy_book), but rather the tagged

[32]https://support.google.com/analytics/answer/1033867?hl=en

URL (https://www.example.com/buy_book?utm_campaign=newsletter-jan-8). This expanded, tagged URL gets tracked in Google Analytics as the visited page (instead of the untagged original), and Google can now parse the information out of these tags. In this case, Analytics knows that the "utm_campaign" was equal to "newsletter-jan-8".

In the remainder of this section, I'll elaborate on the mainstream Google Analytics URL builder approach, but marketers using other analytics platforms can, of course, set their own parameters and program their software to parse these instead.

Google Analytics URL Builder parameters

A few words to Google AdWords users: Be aware that AdWords can be set to automatically fill in these tags for you.[33] Compared to setting the parameters manually, this saves you a lot of time, but you should be aware that automatic tagging constrains your usage of the five URL builder parameters such that they must be consistent with Google's ideas—otherwise you'll create a horrific mess in your Analytics reports.

That leads us to our next point: What is the intended use, as per Google, for each of their five parameters?

Campaign Source: (Required parameter) Represents the platform that sent you traffic (e.g., "google", "facebook", "twitter", or "linkedin"). If you've bought advertising or done deals with smaller websites—be that for banner ads, guest blog posts, or newsletter appearances—write the names of these websites here. Indeed, it's possible that you might be the owner of the campaign source website, as could happen when you promote your new B2B startup on your personal website. When marketing through email campaigns, it's tempting to write "email" as the source (or even the name of the company delivering the newsletter, such as "mailchimp" or "sendgrid"). Don't do this. Instead, write down a descriptive name for which *mailing list* you are promoting yourself on—is it your "pre-launch-list" or your "new-rental-alerts-

[33] https://support.google.com/adwords/answer/1752125?hl=en

www.dbooks.org

subscribers"? By tagging as above, you'll be able to determine how successful each traffic source has been, and you'll be able to apportion future marketing spend accordingly.

Campaign Medium: (Required parameter) Represents the mechanism used to deliver your advertising, such as "paid-advertising", "social-media-post", "banner-ad", "offline", "email", or "affiliate". In many cases, Google Analytics already assigns a campaign medium without your doing anything. You'd be advised to adopt their conventions:

- "direct" (when someone typed your URL into the web browser, or visited a bookmark they previously left to your website, or clicked on a link to your website from an instant messaging app)
- "organic" (traffic from search engines)
- "referral" (traffic from other websites which link to you)
- "cpc" (cost per click)
- "social" (i.e., social networking websites like Twitter, Quora, and Reddit)
- "email"

The point of the campaign medium dimension is to determine which mechanism for delivering marketing campaigns gave you the highest return. Does your Christmas sale campaign perform better through the "email" medium, or through "cpc" or "offline-postering"? Campaign medium is there to help you find this out.

Campaign Name: (Required parameter) Used mostly for distinguishing between efforts that would otherwise be impossible to tell apart (e.g., different email newsletters, different postering campaigns, and, most commonly, different CPC advertisements). Often, the campaign name will mirror that of specific categories of products, leaving taggings like "used-iphones", or "used-macbooks". Within my own campaign name tagging, I like to include information that helps me distinguish between remarketing and non-remarketing advertisements. I do this by appending the abbreviation "rm" onto my remarketing campaign names

(e.g., “used-iphones-rm” vs “used-iphones”).[34]

Campaign Term: (Optional parameter) Called campaign term because Google AdWords auto-tags this with the search term that triggered that particular click, enabling you to better understand what keywords you should advertise on. I would only fill in this tag when you are targeting keywords (or their equivalent, such as likes) on some other advertising platform.

Campaign Content: (Optional parameter) Google AdWords automatically populates this field using the headline of your adverts. Given the ever-expanding variety of advertising headlines you'll experiment with in your quest for high conversion rates, this leads to a semantic mess. You are best off completely ignoring this parameter.

Dealing with grossly oversized URLs

A tagged URL can end up long, cumbersome, and unsightly. Because of its size, it looks unprofessional when shared online, and it can be completely unwieldy when printed out for offline campaigns (e.g., when you want to track the success of your postering campaigns).

To restore concision, you can use a URL-shortening service like Bitly (https://bitly.com/) to package the URL into a bite-sized bundle, ideal for sharing (especially offline, where you cannot expect someone to write out a massive URL that includes tagging parameters).

Tagging consistently

The URL builder is a simple tool that doesn't remember your past preferences for how you tagged URLs given out in previous marketing campaigns. This opens up a breeding ground for inconsistency should URLs be tagged one way today and another way tomorrow. If campaign source is “facebook” today, “Facebook” the day after, then “Face book” the third day, then data

[34] It seems to me more logical to capture this distinction with the campaign medium dimension, but, alas, Google AdWords automatically tags both normal and remarketing campaigns as “cpc”, meaning this more sensible option is out.

www.dbooks.org

which should appear clumped together under the one category will end up unnaturally dispersed across three.

To avoid this pitfall, create a shared spreadsheet that records exactly how your company tagged URLs in previous campaigns. Doing this guarantees reporting consistency to any marketer with the discipline to reference this spreadsheet and the intelligence to write new tags such that they sit harmoniously with the old.

Better yet, this spreadsheet should be programmed with macros that automatically generate the tagged URLs based on parameters the spreadsheet user places in other cells. This process beats giving yourself arthritis by using Google's URL tagger to individually tag each and every URL.

Precautions

Google Analytics data, once written, can never be modified, deleted, or edited. If you slip up, you'll have to live with the consequences of your mistakes forever.

Generally speaking, there are two ways to mess up Analytics data. The first is mistakes in your own website's code (e.g., an on-site bug reports ten hits to Analytics for every one hit on your website). The second category of blunders is mistakes in your Google Analytics account, the most common example being when you create an incorrect Google Analytics filter. (You are probably accustomed to thinking of software filters as windows that apply over underlying data, changing what you see right now but leaving the original intact. But Google Analytics filters are different—they permanently modify or remove data, so getting one of these filters wrong will obliterate your data.)

Early detection of error is the name of the game. Best of all is detection before code is deployed to the live server and permanent damage is done. This could be achieved by having a policy of peer-reviewing recent changes or writing automated tests that probe your Analytics implementation for errors. Next best after

pre-deploy detection systems are damage minimisation systems designed to quickly detect problems as they occur on your live server. This could be achieved by having automatic alerts that get sent whenever Analytics data is out of normal bounds, or by having someone on your team manually monitor the "real time data" report during the minutes after a dangerous change was made to the website.

With regards to the second broad category of errors—those occurring within your Analytics account—you can protect yourself by creating backup profiles. For instance, you could have a separate Google Analytics profile named "Unfiltered Data" that stands back and gathers all your data without the potentially damaging effects of any special filters or other special Analytics features. The purpose of this profile is nothing other than to serve as the ultimate backup and source of truth if all else fails. If you want to be especially careful, you could also create new "testing the waters" Analytics profiles before effecting any potentially damaging change. Once you are comfortable that your planned change won't do any collateral damage, you can switch out the "testing the waters" profile for your business-as-usual primary profile.

Be Disciplined in Making Annotations

Anomalous occurrences in your business's life—such as massive once-off email marketing campaigns, major press coverage, or technical issues that affect tracking—can create massive spikes or sudden gaps in your Analytics data.

In the months or years that follow, these anomalies will confuse you (or the employee taking over your role). Worse yet, if not accounted for in future analyses, these anomalies can contribute to you making false conclusions about your business and thereby making poor business decisions, all the while thinking your error is justified by the data.

These future problems can be mitigated by adding annotations to your Analytics account. These are nothing more than simple text notes, placed at certain dates, whose purpose is to alert future

www.dbooks.org

readers about the anomaly.

Send Alerts

Google Analytics can be configured to email you whenever something remarkable happens. I will suggest four basic alerts that every webmaster can benefit from.

1. Alerts when there is no traffic *at all* over the past 24 hours. At their most dramatic, these alerts inform you about catastrophic infrastructural bugs that take down your entire website. More commonly though, these alerts inform you about bugs in your Google Analytics implementation (e.g., your tracker unexpectedly stopped showing in your HTML template). Knowing about these Analytics problems 24 hours after they occur is important, because it keeps the damage to your data contained. (There is no way to retrospectively repair Analytics data...)
2. Alerts about there being no conversions in the past hours/-days. These would inform you of newly introduced problems that seriously affect the key flows on your live website (e.g., issues with the checkout process preventing purchase conversions). Think of these alerts as basic health tests—assistants to automated software tests, if you will.
3. Alerts about there being a relative reduction in the conversions compared to a previous period (e.g., 30% down). These might inform you that a previous change in marketing strategy was a bad decision, as evidenced in the reduction in sales.
4. Alerts about traffic spikes. These will mobilise you to inquire about what caused the spike (e.g., a mention of your brand on Twitter by a big influencer), how you can capitalise on the buzz, and how quickly you can scale out the servers.

Analytics Blocking

If you're scratching your head about discrepancies between your figures as recorded in Analytics and your figures recorded in

your server's database or logs, then take heed of the following: A small but rapidly growing proportion of internet users are impossible to track with Google Analytics (or similar technologies), owing to their use of tools for opting out of or completely blocking third-party tracking software. In 2016, Quantable found that as much as 11% of traffic is now untrackable, although they noted that the sample they based their calculation on was skewed towards younger, male, internet-savvy web surfers, the demographic most likely to install privacy-protecting software.[35] Regardless of what the exact percentage of untrackable traffic is in your field, analytics blocking is a development to keep your eyes on.

[35]https://www.quantable.com/analytics/how-many-users-block-google-analytics/

www.dbooks.org

13 | Conversion Optimisation

Its Advantages

Conversion optimisation is widely lauded in marketing communities. This praise is justified, as the techniques are every bit as valuable as those who hype them suggest. Here's why:

1. Seemingly small improvements in *absolute* conversion rates lead to amplified *relative* results on revenue (and, subsequently, profits). Suppose that you sell shoes in an ecommerce store. What difference to your business's revenue would increasing the percentage of people who click "Add to Cart" from 1% to 2% make? All things being equal, twice as many visitors would add the product to their carts, so twice as many would buy the products, meaning that you would effectively *double* your revenue (and profits). As such, a "mere" 1% improvement in *absolute* conversion rates, could mean a doubling in *relative* conversion rates, which, when converted to lifestyle benefits, can mean the difference between inhabiting your current apartment and one twice its size.

2. Big efficiency gains in sections of your sales channels where only a tiny subset of your visitors flow only bring minor increases in revenue. By contrast, even modest efficiency gains in components of the conversion funnel that *every customer necessarily passes through* bring enviable improvements in overall revenue. For example, say you generate leads for a product through 100 different off-site advertisements, each of which delivers 1% of your total traffic. Suppose you double the (click-through) conversion rate of *one* of these 100 advertisements. Even taking this doubling into account, you'll only see a 1% improvement in *total* revenue because the other 99% of advertising traffic passes through adverts which are no more effective than they were before. But if you had instead focused your improvements on the on-site landing page, this would have affected *every* visitor; instead of improved

conversion rates for 1% of leads, there would be a boost for 100% of them. As such, a doubling in the landing page effectiveness would also double overall revenue, making this (and other on-site conversion optimisations) an extraordinarily efficient place to focus on improving conversion rates.

3. There is always some finite limit to how many potential customers a business will be able to reach through their paid ads, SEO, email newsletters, partnerships, and other lead generation techniques. For a company like Google, this limit is high, approximating the number of people in the world with internet connections. But for the rest of us—those of us whose businesses are serving much narrower niches—this number is far lower. Some businesses might never be able to reach more than, say, 1,000 leads per week—even when we are doing everything in our power to get our messages out to the public. At this point, we are said to have reached our maximum *advertising reach*. Any further marketing efforts on our behalf are better spent on converting a *larger proportion* of our existing leads into paying customers than on pointlessly spinning our wheels in the doomed search for even more leads. On-site conversion optimisation is uniquely powerful in these situations because it unleashes new revenue growth in businesses that have already exhausted their sources of new leads.

Underrated

Given its wonderful benefits, why then is conversion optimisation so underused?

1. Conversion optimisation—at least traditionally—always required that there was someone on the team who knew how to change the source code and then deploy the modifications. Want to change the name of the "Pricing" link to "Plans and Pricing"? Well, you'll need your engineer's help—and they are already *swamped* with bug fixes.

Nowadays, services like Visual Website Optimiser enable non-coders to carry out simple modifications (like text changes), but

www.dbooks.org

alas, straightforward modifications like these compose but one galaxy in the universe of possible optimisations. More involved modifications (e.g., to the checkout flow or pricing structures) still require complex rewrites of the software, which means that, despite all the recent technological innovations, conversion optimisation can still be expensive to experiment with.

2. Completely separate to the above concerns, there is sometimes the awkward matter of public perception to consider. Were you to lower your product's price from €200 to €120 in order to see if this change optimised total revenue, the customers that bought at €200 would become enraged at learning that their friends had bought the exact same product four hours later for only €120. The fear of facing these problems can scare marketers away from experimenting with conversion optimisation.

3. Conversion optimisation—when done right—asks that its practitioners learn, implement, and patiently apply the statistics of A/B testing. But website entrepreneurs already have their plates full with the millions of other little things they have to do. If your entire website feels like it is crumbling to pieces from the inside, taking time out of your schedule to dive into esoteric statistics often isn't the highest priority.

On-site vs Off-site

The concept of conversion optimisation traditionally applies only to on-site improvements in conversion rates (i.e., optimisations that happen on *your business's website*). Of course, it's possible for conversion optimisation ideas to be applied to the business flows that lie outside your website—for example, to the advertisements that bring new leads there in the first place. We'll distinguish these particular optimisations by referring to them as *off-site* conversion optimisations, as opposed to *on-site* ones. In general, you can assume that whatever we say about on-site conversion optimisation also applies to its off-site cousins.

Beware of Shallowly Applying the Lessons of Others

There is a whole cottage industry of online marketers who blog about how they improved their conversion rates by doing this or that. Large as this group may be, there is an even greater herd of internet marketers reading these blogs and taking their conclusions as infallible gospel. The example of these unthinking readers is not one to follow. It is unwise to blindly apply other people's conversion optimisation tactics because so much depends on your own unique set of circumstances. Instead, you should treat their findings merely as suggestions of experiments to try out and measure, rather than certainties to steal and implement without further thought.

(Of course, you have to launch your business with *something,* so it's a good idea to follow the community conclusions in your version#1—just remember to stay sceptical in subsequent evaluations.)

I would like to demonstrate this danger by telling you about a revealing chain of events. In July 2011, Signal v. Noise blogged that a text-heavy long-form landing page was 37.5% more effective at converting than their previous short-from page.[1] Reviewing these results and wishing to condense the complexities of conversion optimisation into a few easily followed rules, you may be tempted to generalise from Signal v. Noise's experience the rule that "long-form landing pages are better than short-form ones".

But Signal v. Noise's data does not justify so broad a conclusion. All that their experiment showed was that one particular long-form landing page performed better than one particular short-form landing page, for one particular company serving one particular market.

Indeed, events a few weeks later vindicated this scepticism: In August 2011, Signal v. Noise published the results of another test

[1] https://signalvnoise.com/posts/2977-behind-the-scenes-highrise-marketing-site-ab-testing-part-1

www.dbooks.org

to the same landing page.[2] This time, they pitted their original short-form page against an even shorter-form page which had a high-resolution photo of a customer filling up the entire screen. The results: This new design, even though it was extremely short, performed 102.5% better than the original design, trouncing the 37.5% improvement won by the long-form variation. If you had taken the first experiment as evidence for the general superiority of long-form copy, you would be forced to eat humble pie.

The point of all this is that conversion optimisers ought to recognise that their conclusions rest on the shaky foundations of past experience robbed of context. In addition, the verbal explanations given by others for obtained experimental results can sometimes misattribute causal factors, and are therefore also suspect. If a blogger tells you that naming the link to their signup page "Plans & Pricing" brought them a higher conversion rate than naming the same link "Pricing", how can you be sure that the active factor contributing to the conversion lift wasn't something else, like phrase length, number of words, or the presence of an ampersand? Without carrying out your own tests, you can't really know, and it would be rash to blindly accept these verbal explanations. The sober marketer reads other people's successes as evidence of the existence of levers and dials worth experimenting with, not as instructions about the levels and positions these controls should be set to.

And with that we reach our next section.

Areas Worthy of Experimentation

A. Conversion funnel length

Every stage of a conversion funnel has some leakage (40% of traffic seep away here; another 22% dips out there...). Website owners invest untold energies into patching up these leaks so that even more leads reach the end of the funnel and convert.

But no stage of a conversion funnel will ever be *completely* sealed

[2]https://signalvnoise.com/posts/2991-behind-the-scenes-ab-testing-part-3-final

against leakage, and, for this reason, the most efficient possible modification to a funnel is usually to rip out *entire stages*.

Do you *really* need a separate page to collect billing addresses when you deliver your service digitally and can attain the necessary taxation information about that customer from Paypal's API? Is that extra "confirm your order" page *really* necessary, given that (1) the customer has already paused to enter in their credit card details on a prior page and (2) you offer 100% refunds, no questions asked? Is it *really* necessary for the customer to register before purchasing? Sure, this may be standard practice, but perhaps the crowds are mistaken? Jared M. Spool writes in *The $300 Million Button* that a mandatory registration step intended to ease repeat purchases only served to drag down their conversion rates. His usability tests found that customers were "not here to be in a relationship," and that the registration requirement deterred them, causing them to abandon their purchases. The fallout was phenomenal: By removing this registration step, sales increased by 45%, and the retailer increased revenue by $300,000,000 that year.[3]

Amazon's one-click order feature seems to be motivated by similar considerations; it must surely count as the ultimate exercise in minimising funnel length.

Trimming funnel length isn't just about cutting the number of pages in the flow on *your* website; more generally, it's about reducing the amount of tedium and distraction that a customer must trudge through before completing their order. From this point of view, you should consider the following factors when optimising your funnel:

1. Steps behind the scenes

There exist conversion steps that customers pass through even though these steps don't correspond to individual pages on your website.

Perhaps the most commonplace example of this is the confirma-

[3] https://articles.uie.com/three_hund_million_button

www.dbooks.org

tion email click required by some websites to finalise new user account registrations. This step is a 10-ton conversion dampener.

Another (more subtle) dampener would be the very presence of a coupon box on a checkout page. This box alerts the potential customer to the possible existence of discount codes and invites them to leave your website and google for coupon codes, or subscribe to your mailing list and wait until a discount becomes available. In both cases, this distraction might cause you to bleed leads and lose sales.

Yet another conversion dampener in this category is asking a lead for information not normally held in memory—e.g. passport numbers, VAT numbers, scans of documents, or credit card limits. This information demand compels anyone without a reference at hand to drop out of your funnel. If at all possible, ask for such information *after* the sale has been completed.

2. Number of form fields

Funnel length is not just measured by the number of distinct pages; after all, even a one-page funnel can be exhaustingly long. Each additional field that needs to be filled in counts as an extra step in your funnel, regardless of whether this additional field shows up on the same or a different page. Do you *really* need to ask the customer how they heard about your website? Do you *really* need to ask them to upload a document that they probably won't have at hand *right now*?

The more form fields you add within a funnel, the more you risk losing your potential customer. This suggests that, when optimising your funnel, you ought to remove any fields that aren't totally necessary, or least make as many fields as humanly possible into optional ones. You can always ask for this information after the lead has converted into a paying customer.

3. Painfulness of form fields

There are all sorts of things you can do to make your form fields less onerous.

For one, you can display your validation error messages in-line (right next to each invalid field instead of all bunched up together at the top of the screen). In-line validation messages are known to improve conversion rates, and you'd be unwise not to avail yourself of this easy win.[4] Ideally, Javascript-powered in-line validation should occur in real-time as the user fills out each form field, as opposed to only happening after the user completes the entire form and sends the data to your server.

Then there are autofill fields: If someone enters their street name, you should populate their ZIP code without the user having to type anything. Indeed, maybe you can even prevent the need to type in ZIP codes at all, at least when you're able to detect it from their browser's geolocation. As well as all this, you should always autofill any form fields with any existing data you have about that user, saving them the need to type again.

The third measure in reducing painfulness of form fields is making your input validations more forgiving. You really shouldn't be rejecting people's credit card numbers or phone numbers just because there were hyphens or spaces in places you didn't expect. Take a little time to write validation code that's a little more accepting of your customers' input. There's no point in having a bouncer so strict that no one gets into to your club anymore.

4. Steps added by third-party providers

Many websites collaborate with third-party web services which require the customer be sent away to the third party's website at some point during the sales funnel. A common example would be Paypal briefly sending customers off to Paypal.com for credit card payments. These off-site service providers add a number of steps to the sales funnels, so it's important that these extra steps be reckoned when calculating *total* funnel lengths. Thus, when choosing third-party software service providers, carefully consider the number of steps they insert into your conversion funnels. It might well be a better idea to choose a service provider that charges double or even triple as much, so long as they re-

[4] In-line validations brought a 22% conversion lift – http://www.getelastic.com/real-time-inline-validation

www.dbooks.org

move a step of your sales funnel and thereby improve overall conversion rates.

5. Number of choices a customer must make

Professors Iyengar and Lepper found through their experiments that people are more likely to buy products (jams or chocolates, in their case) when presented with a small number of choices (6 items) compared to when presented with more options (24 or 30 items).[5] By extension, stocking excessive numbers of similar products on your website can also hamper conversion rates. I hypothesise that this is because having additional products introduces an implicit extra step in your conversion funnel: the comparison stage. Potential customers who are straining to tell the difference between many almost indistinguishable products may suffer a bout of choice overload that could poison an otherwise sure-fire sale. Remove this comparison stage by stocking less items.

B. Funnel order

Once you've trimmed down your funnel *length*, it's time to start thinking about the *order* in which you display the various retained stages.

Much like difficult family relations, certain stages of your conversion funnel are here to stay—regardless of your feelings towards them. Whether it's the need to ask for credit card details before payment, scans of passports for government-mandated identity verification, usernames and passwords for account creation, or log-in details for a connected Twitter account, these annoying stages often serve to scare potential leads off the path to conversion.

As a website owner, you don't realistically have the option to banish these steps outright, but you can shuffle them about so that they rear their annoying heads only when their corrosive effect on conversion rates is minimised. Just as the waiter presents the restaurant bill only after giving you a free shot of mango

[5]https://faculty.washington.edu/jdb/345/345%20Articles/Iyengar%20%26%20Lepper%20(2000).pdf

schnapps, it may be more prudent to ask your customer for payment details *after* they've scheduled a hairdresser, masseur, or cleaner to arrive this afternoon. Relieved that their desired service is already on its way and afraid of having to go through the whole process again on some other website, the customer will be motivated to endure the pain of entering in credit card details.

How do you decide how to reorder conversion funnel steps? This choice involves two classes of considerations. The first is the effect on conversion rates (to be determined by experimentation). The second class of considerations are those pertaining to technical/programming issues. Allow me to make two general observations as to the flavour of these technical issues:

Observation 1: *The further you stray away from that which is well-worn in the world of web architecture, the more you should expect trouble and expense with implementation.*

Here is one of many examples: Most websites use "batteries-included" libraries for managing their login and logout functionality instead of building their own solutions from scratch. This saves days or even weeks of development time. But these libraries, being standardised, assume typical, everyday website flows (e.g., the user creates a user account before having access to special features or being allowed to pay). Generally speaking, these libraries are unsuitable for alternative website flows (e.g., ones which delay user account creation until the last possible moment). If you wish to keep using the standard account management libraries, you'll need to resort to custom hacks like the creation of zombie user accounts that must later be awoken. But all this customisation causes development time and bug incidences to balloon.

Observation 2: *The more your architecture allows a user to postpone inputting information, the more you should expect problems attributable to this information being missing.*

This is because some proportion of your users will never return to provide this info, causing postponement to drift into chronic

www.dbooks.org

absence. This causes a schism to form between the users with and without this data, which in turn requires that your website's algorithms and administrators be trained to also handle the cases when data X or data Y is missing. What's more, you'll also have to decide how to respond when data which has long been missing finally appears. For example, the authors in my notes publishing platform are paid monthly royalties to their Paypal accounts. But because I allow authors to set up shop without inputting their Paypal email straight away, it sometimes takes authors up to a year to send us their Paypal account details. Once these details are received, our software needs to calculate the royalties accrued up to that date and pay this amount. The software needed to do this requires a fairly substantial amount of code, none of which would be necessary if we had everyone's Paypal account from the get-go. (But, on the flip side, requiring a Paypal account up front would drive away too many potential authors, so I believe that this trade-off is justified. This shows that the bump in conversion rates came at a cost in technical complexity…)

Some software companies have adopted what at first seem like outlandishly long conversion funnels. Central to their strategy is a free demo approach, which involves an inversion of the normal business sequence of paying today and receiving tomorrow. Their funnel sequence goes like this: "Create your account today, enjoy our service for the next 30 days free of charge, then begin paying monthly, but only if you want to continue using us." The power of this "free demo" approach is that the user becomes gradually dependent on the service, having grown accustomed to the interface, trained their employees, integrated the software with the rest of their toolkit, and generated a history of data sitting within the platform. Assuming that the software works reasonably well, inertia all but guarantees that the user won't swap the supplier for another when the bills eventually start arriving. The effort involved in switching becomes too high to justify.

C. Number of clickables per page

Next time you are shopping on Amazon, pay attention to how their website sheds navigation menus as you tunnel deeper down

their conversion funnel. The reasoning behind this practice is that every clickable link, photo, or button is a potential distraction that can potentially divert a lead away from buying something. This distraction happens in the dullest of ways—for example, someone finishes typing their address in the second-to-last stage of the sales funnel and then decides, out of curiosity, to click the link to the "team" page before paying. Once on the team page, this lead reads that the company's CEO used to play in a punk band, so off the lead goes off to search for the CEO's old band on Spotify. Swoosh—the lead has left the site and will probably never return to complete their purchase. There's something almost cartoonlike in these precautions: If we leave a big red button in the middle of this airplane control panel, someone will inevitably succumb to their curiosity and slam it down.

Cautious about avoiding this kind of funnel leakage, conversion optimisation experts advise restricting each page to a single goal—one that is big, totally obvious, and eminently clickable, and that leads inexorably towards conversion. In the same way that a forking river divides its waters, each clickable on a webpage diverts away a stream of potential conversions. This leakage ought to be dammed: Once the customer is heading in the right direction, limit yourself to only including a clickable that brings them closer to the final goal.[6]

Since the dawn of the web, internet marketers have loved pop-ups and overlays. The reason why couldn't be simpler: They work.[7] I mention them here because I would venture to guess that a significant part of the reason why pop-ups are so effective is that they reduce the number of clickables on screen and focus attention on some singular goal, such as handing over your email address for a mailing list.

[6]For example, one company changed from having two buttons to one, improving their conversion rates by 36.9% - https://blog.kissmetrics.com/100-conversion-optimization-case-studies/3/

[7]Brian Dean at Backlinko reluctantly installed a pop-up for collecting mailing list subscribers, only to discover that it *more than doubled* his daily sign-ups. He decided to keep the pop-up around, despite his misgivings – http://www.appsumo.com/sumome-double-conversion-rate/?rf=brws

www.dbooks.org

On a similar note, you should avoid linking to external websites in your conversion funnel. Yes, the glowing testimonial on your product page may have been written by a celebrity with a large Twitter following. And yes, you'd like to emphasise the credibility of the testimonial by directly linking to the original utterance. But doing so may have the unwanted side effect of delivering your customer over to the finely tuned engine of infinite distraction that is Twitter, swiftly sweeping your sale away. (Those preferring a compromise might choose to continue linking to the Twitter account, but programming these links to fire up new browser tabs or windows, thereby keeping the customer focused on completing the funnel.[8])

D. On-page copy

Wording strongly influences conversion rates. Although finding the best variations takes time, the journey can be sped up by focusing on trying out radically different ideas as opposed to slight variations on what currently works best. The retailer selling shoe polish for the business executive shouldn't limit their search for the best product ad after pitting eight mildly different catch phrases about perfectly polished shoes against one another; instead, this retailer should test marketing messages corresponding to radically different themes: business success, confidence and status boost, style, convenience when travelling on airplanes, ease of application, ecological friendliness, or money saved by prolonging a shoe's life.

Calls-to-action are your on-page copy superstars. These are the tiny phrases of imperative-mood language contained within or next to your form buttons and funnel continuation links. Despite their potential to elevate conversion rates, calls-to-action are all too often overlooked by internet marketers who instead leave them set to their ugly defaults. But this is unwise. Asking website members to "join now" conjures up images of a religious cult recruitment drive, and unless you're in the business of sadomasochism, you probably shouldn't be asking customers to "sub-

[8]Guide to programming some links to open in a new window/tab: http://stackoverflow.com/questions/15551779/open-link-in-new-tab

mit". Rather than these uncreative and uncompelling defaults, your calls-to-action should remind the lead of the benefit you promise to provide (e.g., "Remember Everything" (Evernote, a note taking and syncing service), "Start Selling Today" (Square, a credit card processing solution) or "Start Meeting People Today" (OKCupid, a dating website).

I'd like to draw a distinction between copy changes that merely play with words and copy changes that have far-reaching consequences for how you do business. For instance, suppose you modify copy so as to promise the customer "refunds for life". To honour this promise, you'd better have the technical capacity to refund money months or years later, the administrative capacity to update your accounts (otherwise you may be taxed on money you didn't make), and the website features to prevent further delivery of your service to a refunded customer. Even though your copywriter only needed to add the three words "refunds for life" to the webpage, the assumptions behind these words could require months of changes tallied in terms of new features, modified workflows, and retrained staff. Not all words are cheap.

E. Images

It won't exactly be ground breaking news to you that images are one of the most important elements to experiment with when optimising for conversions.

As a general rule, the more immersive the image, the better. Images with a 360-degree spin feature were found by DueMaternity.com to give a conversion lift of 27%.[9] Larger product images were found by car manufacturer Hyundai to improve conversions by 62% compared to thumbnails.[10] On a similar basis, having an image zoom feature is likely to improve conversion rates, as does having additional images depicting your products from various angles or used in differing contexts.

Product images should not be reserved for the product sales page

[9] https://www.internetretailer.com/2011/05/26/rotating-images-generate-higher-conversion-duematernitycom

[10] https://vwo.com/blog/multivariate-testing-case-study/

www.dbooks.org

alone. When online retailer Brickhouse Security added mini product images to their drop-down on-site search, they boosted their conversion rates (of customers who searched) by no less than 100%.[11]

The choice of image matters too, especially since the usual raison d'être of an image is to transmit a carefully chosen marketing message. Web hosting company Hawk Host was able to more than double their conversion rate by swapping a photo of a globe for a photo of a giant padlock. This change helped them to visually emphasise their company's focus on security.[12] Another example comes from Apple: They advertise their laptops with photos depicting them being slid into envelopes, showing us just how tiny their computers are.

Another effective use of image is to demonstrate how a product looks while being put to use. This helps the potential customer imagine how the product being advertised could deliver to them the benefit they desire.

Generally speaking, photos of human beings tend to convert better than photos of things. One company improved their conversion rates (of people who contacted their customer support) by 21% just by adding a photo of their customer support representative.[13] (Actually, this company tried out photos of two different representatives; one of them only added 10% to their conversion rates. This goes to show that you should vary *which* people you display on your pages.) Basecamp (formerly 37signals) improved their conversion rates by 100% when they redesigned their home page to focus on a picture of a customer instead of information about their software products.[14]

Another advantage of using photos of people in your copywriting is that their facial expressions and postures communicate emo-

[11] https://www.internetretailer.com/2010/04/29/product-images-site-search-window-boosts-conversions

[12] http://blog.hawkhost.com/2010/02/21/multivariate-testing-a-real-life-example/

[13] http://www.entrepreneur.com/article/220252

[14] https://signalvnoise.com/posts/2991-behind-the-scenes-ab-testing-part-3-final

tions. This can help you to conjure up feelings of comfort, happiness, and more in the hearts of your leads.

One final piece of advice: The pictures on your website should not be stock photos. Even the most amateur pic you yourself take is likely to substantially outperform stocks. An experiment run by a moving and storage company found that a crappy image of their truck beat by 45.05% a stock image of an attractive couple lifting a moving box.[15] (A comparably amateur image of their team of workers wearing their blue overalls performed marginally better, giving a conversion lift of 45.45% over the stock image.)

F. Pricing

The most successful conversion optimisation I ever instituted in my business came from a change that took all of half an hour to implement: I raised the prices by 20%. As I did this, I feared that enraged customers would throw a virtual brick through my shop front, refuse to pay the increased price, and decry our greed in every internet forum there is. In fact, nothing of the sort happened. The number of sales stayed exactly the same, causing profits to increase by the 20% that I had raised the prices by. I won't deny that occasional customers (say five per year) contact us to say that our prices are too high. But these protestations aside, our prices don't seem to inhibit actual buying behaviour. Customers continue to convert at practically the same rate at our higher price point. Indeed, some of the very same people who complained about our price increase nevertheless went on to purchase our products.

I cannot stress how big a difference raising your prices can make. This one tiny change might dwarf all the gains you achieve through every other business activity in the next year.

Increasing the amount you charge for a product is just one way of tweaking pricing. But there are other, more creative ways, to increase profitability by experimenting with new billing mod-

[15] https://vwo.com/blog/stock-photos-reduce-conversions/

www.dbooks.org

els. An acquaintance running a digital mapping service for cyclists used to charge his customers a monthly fee for access to *all* the maps in his system. He dramatically increased revenue by switching his billing model such that customers would instead buy maps of smaller regions. The customers would then "own" these maps forever instead of paying for them with a monthly fee. The typical customer likes the idea of buying to keep, and, moreover, most people have a deeply ingrained instinct to complete collections, with the consequence for my acquaintance that many of his customers bought *every* map in his system, regardless of whether they would ever use them.

All things equal, money now is better than the same money later. For that reason, it's good business to give discounts to customers buying a year's worth of service in advance (as compared to customers buying your service one month at a time). When you bill monthly, the customer may change their mind at any point and switch to a competitor. Had you billed yearly, the customer might have stuck around long enough to become so invested in and hooked on your product that they would no longer consider switching once the next billing cycle comes around. In addition, revenue booked now—as opposed to later—can be used to build new features or to woo new investors. Therefore, it's worth experimenting with extended billing cycles as part of your conversion optimisation efforts.

I'm a fan of bundling, which is when you package a bunch of products together into a discounted collection. Although the margin on each bundled item drops because of the discounting, the average sale size increases thanks to customers now buying three items instead of just one. When done right, this increase in total revenue should more than offset the reduced margin, spelling increased overall profits. If you want to really boost bundle sales, try making the constituent products massively overpriced when bought separately. Although this will sabotage sales of the unbundled products, it will boost sales of the bundle by making it seem like comparatively great value.

Businesses which have corporations or public entities as part of their customer base may consider pricing schemes which capture exponential value from these large and wealthy entities. The typical software service provider has a linear pricing system, one that scales with usage (such as "$5 for each team member" or "$20 per server per month"). These service providers would therefore charge a corporation with 1000 team members the same amount of money as they would 1000 personal team-of-one customers. By contrast, a service provider billing with an exponential value pricing scheme would charge corporations exponentially more by creating special "enterprise" plans. These plans contain features that large corporations *must* have, be that for regulatory reasons, security reasons, tax efficiency reasons, data ownership reasons, branding reasons (e.g., white-labelling), or integration assistance reasons. Under such a scheme, the software provider would instead charge the corporation with 1000 team members five times as much as they would charge 1000 individuals. The higher price charged to corporations could even be used to subsidise the costs of individual accounts.

Sometimes, changing your price tag or your billing model is too daunting a task to undertake. In these situations, there's always the possibility of opting for a merely cosmetic change, such as reframing the current price. Service providers who bill monthly could deemphasise their month charge (e.g., £30 per month) and instead draw attention to their lower amortised daily cost (e.g., £1 per day) or the average price per item (e.g., £0.33 per episode watched, assuming the average user streams three shows per night).

So far, we've shown how you can dress up your price in more attractive *monetary* terms; there's also the possibility to frame your price in terms of non-monetary benefits, such as the number of hours of work saved or the number of server downtime incidents avoided.

www.dbooks.org

G. Timing

Sprinkled through your codebase and business processes are variables determining durations of time. These represent things like how long your software waits before sending a discount to people who abandoned their carts, how many seconds it waits before displaying a pop-up to website visitors, how many days it gives away as part of a free demo before throwing up a pay-wall, how many times per month it emails the people on your mailing list, or on what date it switches on the holiday advertising campaigns. The online marketer concerned with conversion optimisation should be prepared to tweak these timings.

Time influences conversion optimisation in another respect: Sluggish website response times stifle conversion rates. Thus, a "quick win" comes to those who monitor their website logs (or New Relic/Google Analytics reports) for abnormally slow pages (e.g., those with response times > 200ms) and focus on making these pages snappy again.

H. Salience (includes graphic design)

Just because you, the author of your website, are intimately aware of every word on your website, doesn't mean that a random visitor will be. The average person probably spends mere seconds on your website before they move on. It is the job of your visual design to hook passers-by into your world before these precious seconds are up.

The visual design of copy elements—such as their relative size, placement, spacing, or colour—matters. A headline in font size 48 at the top of the screen is going to receive dramatically more attention than your 6th bullet-point out of 10, rendered in font size 11 and submerged deep below the fold. This implies that moving different messages into and out of your design's visual spotlight will affect your conversion rates. Rotating your existing pieces of copy to new positions on your website can potentially be as effective in improving conversion rates as coming up with brand new elements.

As well as an element's font size or positioning above the webpage fold, there are various other visual cues and tricks that a marketer can employ in drawing attention. One is simple visual contrast. If you happen to have one paragraph of text on your webpage contained within a box with a *coloured* background, visitors will be more likely to see it and therefore to read its text, if only for the simple reason that it contrasts with all the other boxes on the page which are set against white backgrounds.

Buttons that attract attention are also more likely to be clicked on. You can engineer a button to stand out by choosing a colour for it that contrasts with the rest of the page, by making it oddly shaped or ridiculously oversized, or by programming it to shake ever so subtly, thanks to a titbit of Javascript.[16]

In fact, this brings us to a more general point: Movement is a surefire route to salience. One example of this idea in practice is how some e-commerce websites pop up a customer service representative chat bubble after 20 seconds of browsing. Another example would be "Please subscribe to our mailing list" bars that appear over a website's header or footer after you move the mouse in a certain way. Yet another example is the updating of an on-site notification counter to indicate that a new message has arrived, as in Facebook or Twitter. Another example again would be the sudden appearance of a breaking news story at the top of a newspaper website's homepage. In all these cases, the sudden change to the website visuals steals away your attention.

Salience doesn't just happen by spotlighting the element in question; it can also happen by *de*-emphasizing non-important aspects of the page. That's why some pages grey out their backgrounds when they want you to focus on converting.[17]

Another way to draw attention to your desired goals is by using di-

[16] Read this page until you notice the jiggling 'Download' button: http://ghostinfluence.com/reddit-traffic-explosion-deconstructing-the-framework-and-tapping-its-potential/

[17] For an example of graying out the background, scroll around this page for about 30 seconds: http://kaiserthesage.com/email-list-building/

www.dbooks.org

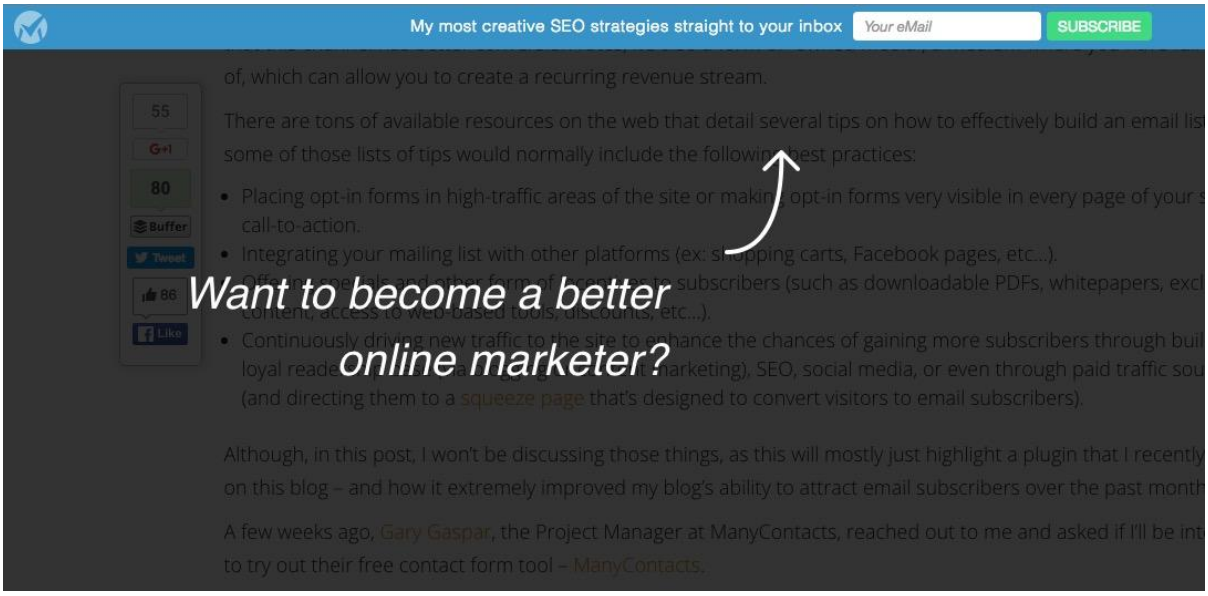


Figure 13.1 – *Source: Kaiser The Sage*

rectional cues, such as the cheesy-looking hand-drawn arrow depicted above. Classier marketers can achieve this same effect with subtlety, for example by inserting a picture of someone whose eyes are gazing directly at an important button, or by having a picture of an animal (e.g., a squirrel or a lizard) whose head is pointing towards the critical sales message.

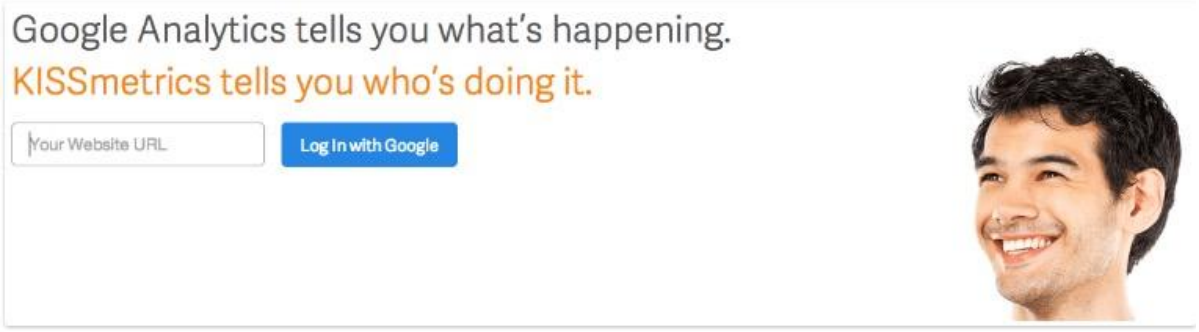


Figure 13.2 – *Source: Kissmetrics*

I. Choice architecture

Owing to our human proneness to inertia and herd mentality, most of us tend to stick with whatever is presented to us as the default option. Professor of Psychology and Behavioural Economics Dan Ariely observed that in countries where being an organ donor is the default, the overwhelming majority of citizens

"sign up". For instance, 99.9% of Austrians are donors. By contrast, in countries where citizens have to explicitly opt in to organ donation, donor levels are dramatically lower: Only 12% of Germans are donors, despite their being culturally similar to Austrians.[18]

The study of how outcomes can be influenced by the way choices are presented is known as choice architecture. And, as you can guess, this field can be applied to conversion optimisation.

We've all seen those checkboxes at the bottom of web forms that read, "I wish to receive periodic updates about new offers". As a webmaster, you can decide whether you leave this box ticked or unticked by default. Amazingly, only 18% of users ever change this default.[19]

There are other ways to leverage this idea in your business. You might suggest one subscription package as a default over the others; you might remind customers that it is customary to buy product X with product Y; or you might go full evil and include some opt-out "extra" (e.g., insurance) in the cart without the customer explicitly adding it. (See low-cost airlines for more on this shady practice.)

J. Available features

I feel a little cheeky suggesting new features as a conversion optimisation tactic. Conversion optimisation is supposed to be about quick wins, whereas building new features involves actual work. But even taking this effort into account, the jump in sales that could follow the addition of a feature highly valued by your customer base could make all the toil in the world worthwhile.

There is a low-cost strategy for web businesses unsure about investing in features of unknown contribution to future conversion rates. Instead of actually building the feature, these businesses could merely promise it within their website copy and observe how this promise influences conversion rates through

[18] http://danariely.com/2008/05/05/3-main-lessons-of-psychology/
[19] http://www.inc.com/articles/2002/10/24718.html

www.dbooks.org

A/B testing. If these tests demonstrate that the promised feature improves conversion rates, then the business could confidently proceed to build it. (Any customers promised the non-existent feature could be told before the payment step that the feature doesn't currently exist, or they could be offered a refund.)

Deciding what new feature to promise needn't require any creative genius—a change as banal as adding an additional payment option can suffice for a revenue gain that would be the envy of many companies. For example, some customers would never pay with Paypal because of a lack of support in their country or its unattractive currency exchange rates. At the same time, other customers, calmed by Paypal's internationally recognised brand name and many years in business, would not trust any provider *except* Paypal. As such, a change as simple as wiring up another payment option can cure many a conversion optimisation ill. A/B tests enable you to discover these efficiencies without the hassle of actually building them.

14 | Statistical (In)Significance

(Warning to mathematicians: I have grossly oversimplified some concepts for the sake of an easier read. My goal with this section isn't technical rigour; rather, it's to give marketer-readers a rough-and-ready intuitive understanding of the main ideas.)

The Theory

In marketing, as in other fields, we observe patterns about what works and what doesn't, and from this experiential basis we form the conclusions that guide our future actions. As an entreprenerd, your thirst for reliable conclusions will be particularly keen when it comes to online advertising, for mistakes here can quickly drain your company coffers. But despite the pressing need for guidance, it would be unwise to form conclusions in haste, lest your reasoning become tainted with a bogeyman known as "statistical insignificance".

The human brain contains a design flaw: It is all too quick to jump to conclusions, even when it's relying on the most paltry and woefully insufficient data. The internet marketer who draws conclusions about advertising performance on the back of thin data is no different from the gambler who flips a coin three times, sees "heads" land thrice, then pronounces the coin rigged to land "heads" *every time*. It would be terribly unwise for this gambler to bet his house on this conviction. After all, it's plenty possible that a perfectly fair coin could land "heads" three times in a row—such is the nature of randomness.

Let's move this discussion back to advertising. Suppose that we have two ads which differed only in their choice of photo. After leaving these ads running (simultaneously) for a week, we get the following results:

www.dbooks.org

Table 14.1 – Our experiment

Photo Variant	Conversion Rate
Cute Kitten	0%
Angry Puppy	9.5%

These results show that the "Angry Puppy" photo variant was infinitely better at converting than the "Cute Kitten" one. Does this mean that it's time for us to remove all our kitten photos and replace them with ones of puppies?

Before we can answer that question, a co-worker barges into our office brandishing the results of her latest advertising experiment—an experiment which happened to use the *exact same* two ads as our experiment did. Her results, however, differed from ours:

Table 14.2 – Our coworker's experiment

Photo Variant	Conversion Rate
Cute Kitten	10%
Angry Puppy	9.7%

Based on her data, she concludes that the cute kitten adverts were ever so slightly *better* than the angry puppy ones. Therefore, it is the puppy ads that ought to be culled! In other words, she's recommending the opposite course of action to what we would. In the face of these conflicting results, how is our marketing department supposed to make any long-term decisions about what to do?

Mistakes have serious consequences here: A misguided marketing conclusion can wreak havoc on our businesses, since we accept a falsity as an axiomatic premise and build our advertising

strategy upon it. We will be unlikely to doubt or revisit these erroneous precepts, because they were based on "facts" we "empirically" observed with our own two eyes. George Bernard Shaw says it best, "Beware of false knowledge; it is more dangerous than ignorance".

What *concrete* effects could all this have? An advertiser, seeing an initial lead in one campaign, might prematurely divert money away from other campaigns that would otherwise eventually prove to be far more effective. This advertiser might throw money at a suboptimal campaign instead of at its far better-performing, albeit later-blooming cousin, and in that same rash moment they also throw away their profits. But that's not all: When writing *future* advertisement copy–be that for CPC adverts, for on-site landing pages, or for printed brochures–this same advertiser will dissect the elements that they believed made the initial advertisement such a "success" and proceed to imbue and pollute the next wave of advertising copy with these same suboptimalities, amplifying and immortalising them.

Let's dig a little deeper by viewing more columns of the results data from *our* experiment (i.e., the first one):

Table 14.3 – Our experiment, with extra data

Photo Variant	Impressions	Conversions	Conversion Rate
Cute Kitten	10	0	0%
Angry Puppy	1,000	95	9.5%

Well it's a good thing we peeked at those extra columns: Now we notice that there was a staggering difference in the number of times each of the two competing ad variants was shown. The 10 impressions of the kitten ad were pitifully few when compared to the 1,000 of the puppy variant. This makes our results for the kitten variant seem far less robust, far less reliable…

Let's see those same extra columns for our coworker's experi-

www.dbooks.org

ment:

Table 14.4 – Our coworker's experiment, with extra data

Photo Variant	Impressions	Conversions	Conversion Rate
Cute Kitten	1,100	110	10%
Angry Puppy	1,000	97	9.7%

Our co-worker's experiment had a larger and more balanced number of impressions in each arm of the trial, as we can see from the table above, and this makes her results feel somehow more solid than ours.

After seeing this expanded dataset, our intuitive general feeling is that we should have less confidence in the results of an experiment when they are based on only a paltry number of observations (or when *one arm* of an experiment has only a paltry number of observations). To give a simpler, more immediate example of this intuition in action, imagine you saw a book rated an "average" of 5 stars on Amazon. If this average were based on the review of only a single reader, you would hardly think this book better than another which was rated lower (say 4.2 stars), albeit on the back of hundreds of reviews.

This shows us that when evaluating results produced from only a low number of observations, we should ask whether what was reported—the "5-star average review" or the "0% conversion rate"—was not just some artefact of random chance, some one-off fluke.

This intuitive idea has not gone unnoticed by mathematicians, who have packaged it up with a delightfully self-explanatory name: "The Law of Large Numbers". Think back to the gambler who bets his house on a coin being rigged. We would pity his stupidity if he bet after seeing a coin land "heads" only three times in a row. But few would doubt the merits of his decision to bet if that coin had instead landed "heads" a million times

in a row. Intuitively—and indeed mathematically—this is a far sounder bet.

But notice that he is still betting. The gambler can never be fully, totally, and absolutely *certain* the coin is rigged. Even after seeing a million "heads" in a row, there is still an infinitesimal yet nevertheless existent chance that a fair coin *could* have given the result of "a million heads in a row". But practically speaking, this is exceedingly unlikely, so he shouldn't let this tiny shard of uncertainty deter him from making a sound bet.

Now we have two extremes: flipping a coin three times, after which it is too early to make a confident bet; and flipping a coin a million times, after which we would feel exceedingly secure putting money down. But what if we don't have the time to wait for a million flips, but still want to bet in confidence? Translated into the business context, what if we don't want to continue our puppies vs kittens experiment for 10 years before deciding which photo was better? By waiting that long, we would have wasted 10 years in showing a proportion of our customers a photo that was comparatively ineffective at effecting sales. Had we figured out which photo was better earlier, we could have had our best foot forward for a much longer time and earned higher profits in the interim.

There is a trade-off between certainty and experiment length. We can intuitively see this by considering how our feelings of confidence would develop after an increasingly long series of coin flips. After 1 flip of "heads", none of us would suspect the coin of being rigged. After 5 flips, we'd start entertaining the thought. After 10 flips, most of us would suspect it, but not enough to place a major bet. After 100 flips, little doubt could remain in our minds, and we'd feel confident about making a wager. After 1,000 flips, we'd be screaming at the bookie to take our money.

The more consecutive "heads" flips we witnessed landing, the more certain we'd feel about the coin being rigged. But given that we're not immortal and that we will never reach 100% certainty with anything in our lives, we must all choose a pragmatic

www.dbooks.org

point where our uncertainty reaches a tolerably low level, a point where we put our hands up and say, "I've seen enough—let's do this thing".

This trade-off point is quantified by statisticians with a figure they dub the *p-value*. Very roughly speaking, the p-value corresponds to the chance you have of being wrong about your conclusion.[1] The p-value is thus a preference, one that represents your desired trade-off between certainty and experiment length. Typically, marketers set their p-value to .05, which corresponds to having a 1 in 20 chance of being wrong. If you're risk averse about making mistakes, you could set your p-value to .01, which would mean you have only a 1 in 100 chance of being wrong (but your experiment would take much longer to reach this heightened level of certainty).

Now we arrive at our last piece of jargon, the idea of *statistical significance*. Broadly speaking, we call experimental results *statistically significant* whenever they have reached our chosen threshold of certainty (e.g., a p-value of .05). Whether or not results reach this point depends on various factors, such as the number of observations and how big a difference in conversion rates we expect to detect.

There is a point we haven't brought up yet, but which is nevertheless crucial to emphasise: Statistical significance requires that our samples be randomised such that they fairly represent the underlying reality. For example, a US presidential election poll administered exclusively to Fox News viewers would over-represent Republican support, since this audience contains a higher proportion of Republicans than America as a whole. The experimenter should have polled a broader section of American society (e.g., by polling viewers of *all* TV channels or by polling random people listed in the telephone directory.)

Back in the world of advertising, an experiment comparing two adverts would be invalid and unrepresentative if the variants

[1] I told you I'd butcher the definitions.

were shown exclusively to old people, even though the product was intended to be advertised to people of *all ages*. Each age group might have a differing likelihood of responding to a particular advertisement, so the results taken from advertising exclusively to old people won't give a representative picture of advertising to everybody.

Let's wrap up the theory by pointing out that by systematically doing statistical significance testing on our marketing inventory, we achieve two desirable ends: We restrain ourselves from making premature bets that might later turn out to be wrong, and we protect ourselves from accumulating false assumptions about what works in our domain. Statistical significance is a stamp of confidence and quality on our results.

The typical advertiser does not test for statistical significance. He takes out a ream of online advertisements, schedules them to run one after another, waits until the campaigns finish, compares the results of each ad in terms of conversion rates, then crowns whichever had the highest rate as the winner to be reused in all future campaigns. But his conclusions may well be unjustified by the data; this advertiser is little more than a snake oil salesmen peddling untruths to himself. They have, as the book title by Nassim Nicholas Taleb goes, been fooled by randomness.

The Practice

Let's switch our attention away from theory and back to the day-to-day problems facing you as an online marketer. The goal of the remaining sections will be to further sharpen your intuition about statistical significance so that you can better apply the concepts to your work.

Know what counts as a sufficiently "large number" in the "Law of Large Numbers"

Imagine we observed the following advertising results:

www.dbooks.org

Photo Variation	Impressions	Conversions	Conversion Rate
Miniature Pony	1,000	1	0.1%
Miniature Goat	1,000	3	0.3%

Upon eyeballing the data, we see that the goat variant tripled its equestrian competitor's conversion rate. What's more, we see that there was a large number of impressions (1,000) in each arm of the experiment. Is this enough to satisfy the aforementioned "Law of Large Numbers" and give us the certainty we need? Surely this means that the "Miniature Goat" is the better photo in a statistically significant way?

Not quite. Without going too deep into the math, these results fail to reach statistical significance (where p=.05). If we concluded that the goat was the better photo, we would have a 1 in 6 chance of being wrong. Our failure to reach statistical significance despite the large number of *impressions* shows us that impressions alone are insufficient in our quest for certain results. This might surprise you. After all, if you saw a coin land "heads" 1,000 times in a row, you'd feel *damn* confident that it was rigged. The math of statistical significance backs this feeling—your chances of being wrong in calling this coin rigged are about 1 in 1,000,000,000,000,000,000,000,000,000,000... (etc.)

So why is it that the coin was statistically significant after 1,000 flips but the advert wasn't after 1,000 impressions? What explains this difference?

Before answering this question, I'd like to bring up a scary example that you're probably already familiar with from the news: Does the use of a mobile phone increase the risk of brain tumours? This is a fiendishly difficult question for researchers to answer because the incidence of brain tumours in the general population is (mercifully) tiny to start off with (less than .01%). This means that experiments need to include huge numbers of people in order to detect a modestly increased cancer risk (e.g., to

detect that mobile phones double the tumour incidence to .02%).

Suppose that we are cancer researchers. If our experiment only sampled 100 or even 1,000 people, then both the mobile-using and the non-mobile-using groups would probably contain 0 incidences of brain tumours. These sample sizes are too small to give us any information. Now suppose that we sampled 10,000 mobile phone users and 10,000 non-users (good luck finding those...) At the end of this experiment, we might count two cases of brain cancer in the phone-using group and one case in the other. A simpleton's reading of these results would be that the incidence of cancer (or the "morbid conversion rate") with phone users is double that of non-phone users. But intuitively this feels like too rash a conclusion—after all, it's not that difficult to imagine that the additional tumour victim in the phone-using group turned up there merely by random chance. (And indeed, the math backs this up: this result is not statistically significant at p=.05; we'd have to increase the sample size about 8 times before we could detect this difference.)

Let's return to our coin-flipping example. Here we only considered two outcomes—that the coin was either fair (50% of the time it lands "heads") or fully biased to "heads" (100% "heads"). Phrasing the same possibilities in terms of conversion rates (where "heads" counts as a conversion), the fair coin has a 50% conversion rate whereas the biased coin has a 100% conversion rate. The absolute *difference* between these two conversion rates is 50% (100% - 50% = 50%). That's stonking huge! For comparison's sake, the (reported) difference between our goat and pony photo variants was only .2%, and the suspected raised cancer risk for mobile phone users was .01%.

Now we get to the point: It's easier to detect large differences in conversion rates. They display statistical significance "early" (i.e., after fewer flips or fewer impressions, or in studies with smaller sample sizes). To see why, imagine an alternative experiment where we tested a fair coin against one ever so slightly biased to "heads" (e.g., one that lands "heads" 51% of the time). This

www.dbooks.org

would require many many coin flips before we would notice the slight tendency towards heads. After 100 flips we would expect to see 50 "heads" with a fair coin and 51 "heads" with the rigged one, but that could easily happen by random chance. We'd need about 15,000 flips to detect this difference in conversion rates. By contrast, imagine detecting the difference between a coin biased 0% to "heads" (i.e., always lands "tails") and one biased 100% to "heads" (i.e., a 100% difference in conversion rates). After 10 coin flips we would notice that the results were either all heads or all tails. Would there really be much point in continuing to flip 90 more times?

This brings us to our next point, which is really just a corollary of the above: Small differences in conversion rates are near impossible to detect. The easiest way to understand this point is to consider what happens when you compare the results of two variants with identical conversion rates: After a thousand, a million, or even a trillion impressions, you still won't detect a difference in conversion rates, for the simple reason that there isn't one!

Bradd Libby, of Search Engine Land, calculated the rough number of impressions necessary in each arm of an experiment to reach statistical significance.[2] He then reran this calculation for various different click-through rate (CTR) differences, showing that the smaller the expected conversion rate difference, the harder it is to detect.

Experiment	CTR of Control	CTR of New Variant	Impressions Needed
No. 1	2%	14%	100
No. 2	2%	3%	4,000
No. 3	2%	2.05%	1,000,000
No. 4	2%	2%	Infinite

Notice how in the final row an infinite number of impressions are

[2]http://searchengineland.com/statistical-significance-not-just-for-geeks-anymore-38105

needed; you will never detect a difference, because there is none to detect! The consequence of this is that it's not worth your time pursuing tiny expected gains; instead, you'd be better off going for a big win that you have a chance of *actually* noticing.

The difference between ultimate and intermediary goals.

As a customer moves along your conversion funnel towards a purchase, they pass through proximate waypoints, intermediary links in the daisy chain that turns Facebook Ads traffic into website visits and website visits into sales. These stepping-stone goals comprise things like clicking on online advertisements, signing up for newsletters, pressing "add to cart" or "pay now" buttons, and so on and so forth. These waypoints are omens of good things to come, but they are not themselves the ultimate destination.

In general, it's easier to measure the achievement of intermediary goals (like advert click-through rates) than to calculate how these little victories contribute towards the eventual goal that matters most, which is usually sales made. But taking the easier path and optimising for intermediary goals brings dangers. You might only reach *local maximums* as opposed to *global maximums*. To visually understand the difference between these two concepts, imagine that you are trekking through a mountain range. It might appear that the peak before your eyes is the highest in the range, but this appearance would be deceptive if behind that mountain were an even higher one hidden from your view.

Leaving the mountains behind and returning to the drabber world of online advertising, it's possible that your Facebook advertisement with the highest click-through rate is the worst performer in your entire inventory when you measure its success not in terms of click-through rates, but rather in terms of that ad's contribution to eventual sales. Maybe the ad promised Facebook members a free sample, and lots of people clicked, got their freebies, but then never ordered anything that would bring you revenue. By contrast, a different ad, which might only have a modest click-through rate on Facebook, might have attracted big

www.dbooks.org

spenders who would have gone on to do wonders for your profits. If you optimised for click-through rates, you would have won the battle but lost the war.

For these reasons, you should focus your statistical analysis on completions of *ultimate goals* rather than on completions of intermediary ones. In other words, you should ascribe "heads" or "success" to goals like sales rather than to ones like advertising clicks. Of course, to be able to pull this off, you'll need tracking capabilities that can reach across the entire sales funnel, along with more traffic that you'd otherwise need for less sophisticated calculations. For example, if 1,000 ad impressions lead to 95 clicks, and 95 clicks lead to 5 sales, then by switching from measuring success in terms of clicks to measuring in terms of sales, you are left with a much lower number of successes in your analysis—5 instead of 95. This causes the expected conversion rates to drop from 9.5% to .5%, which, as we mentioned in the previous section, makes it harder to reach and detect statistical significance.

This last requirement of high traffic is going to be problematic if your business is just starting out; there will probably be insufficient data for you to arrive at statistically significant conclusions about ultimate goals within a tolerable timeframe. In these cases, it's expedient to lower your standards and ascribe success to intermediate goals, accepting the consequent risk that you may only be optimising for a local instead of a global maximum.

Understand the scope and business value of your results.

Look at the following table of data:

Photo Variation	Impressions	Conversions	Conversion Rate
60s hippy (incumbent)	1000	140	14%
80s rocker (challenger)	1000	300	30%

After applying a statistical significance test, we would see that the 80s rocker photo outperforms the 60s hippy photo in a statis-

tically significant way. (The details aren't relevant here.) But we need to be careful about *what exactly* this statistical significance implies, lest we get carried away.

Our first instinct upon seeing the above data would be to interpret it as proving that the 80s rocker photo converted at a 16% higher rate than the 60s hippy photo, where 16% is the difference by subtraction between the two conversion rates (30% - 14% = 16%). But calculating the conversion rate difference as an absolute change (rather than a relative change) would lead us to *understate* the magnitude of the improvement. In fact, if your business had these results, you would more than *double* your conversions (and, all things being equal, your revenue) by switching from the incumbent 60s hippy pic to the new 80s rocker pic. (Specifically, you would have a 114% improvement, calculated by dividing the improvement, 16%, by the old rate, 14%.) Because relative changes in conversion rates are what matter most to our businesses, we should convert absolute changes to relative ones, then seek out the optimisations that provide the greatest improvements in these terms.

But before we get too excited about the huge relative improvement we'd enjoy by switching to the 80s rocker pic, there is a sobering point to remember: The usual statistical significance tests only go so far as to certify that one variant *performs better* than the other; they guarantee nothing about the *extent* of the improvement. Applying this principle to our current example, our 95%+ confidence that the 80s rocker picture is *better* doesn't give us the same kind of certainty as to it being *114% better*. Without further testing, assisted by more complicated math, we can't confidently say *how much better* the rocker variant really was than its competitor. This is an important point to remember when interpreting what the tests say and what they do not.

There is a distinction between *statistical* significance and *practical* significance. Some observations stamped with statistical significance may be of no consequence or value to our businesses. Suppose that you discover that some advert variant gives you a

www.dbooks.org

relative conversion rate boost of .01%. Even if this result is *statistically* significant, the .01% conversion rate boost is unlikely to be worth the accompanying cost and effort in modifying your advertising accounts and landing pages. It is statistically significant but practically insignificant.

Similarly, a 400% conversion rate improvement would be of little value if it only applied when advertising to a tiny niche, one that will be exhausted after only a week.

We'll close this section by taking a quick look at what it means when some result is statistically *in*significant. I've reprinted an example from earlier which had a p-value of approx .15, making it statistically insignificant.

Photo Variation	Impressions	Conversions	Conversion Rate
Mini Pony (incumbent)	1,000	1	0.1%
Mini Goat (challenger)	1,000	3	0.3%

Does the lack of statistical significance imply that there is a full reversal of what we have observed? In other words, does the statistical insignificance mean that the "Miniature Pony" variant is, despite its lower recorded conversion rate, actually better at converting than the "Miniature Goat" variant?

No, it does not—not in any way.

All that statistical insignificance says here is that we *cannot be confident* that the goat variant is better than the pony one. In fact, our best guess is that the goat *is* better. Based on the data so far, there is an approximately 85% chance that this is true (1 - the p-value, .15 = .85). The only issue is that we cannot be confident to the degree dictated by our chosen p-value—to the minimum level of certainty we wanted to have.

One way to intuitively understand this idea is to think of any recorded conversion rate as having its own margin of error. The

pony variant was recorded as having a .1% conversion rate in our experiment, but its confidence interval might be (made-up figures) .06% above or below this recorded rate (i.e., the true conversion rate value would be between .04% and .16%). Similarly, the confidence interval of the goat variant might be .15% above or below the recorded .3% (i.e., the true value would be between .15% and .45%). With these margins of error, there exists the possibility that the pony's true conversion rate would be at the high end (.16%) of its margin of error, whereas the goat's true conversion rate would lie at its low end (.15%). This would cause a reversal in our conclusions, with the pony outperforming the goat. But in order for this reversal to happen, we had to take the most extreme possible values with our margins of error—in opposite directions to boot. In reality, these extreme values would be fairly unlikely to turn up.

Run your tests in parallel, not consecutively.

If you ever catch yourself comparing the results of two advertising campaigns that ran one after the other (e.g., on consecutive days/weeks/months), then stop right now. This is a really *really* bad idea, one that will drain every last ounce of statistical validity from your analysis. This is because your experiments are no longer randomly sampling, a necessity we mentioned in the theory section.

Imagine you are a gift card retailer and you observed that 4,000% more people bought Christmas cards the week before Christmas compared to the week after. You would be a fool if you concluded that the dramatic difference in conversion rates between these two periods was because the dog photo you advertised with during the week preceding Christmas was 40 times better at converting than the cat photo used the following week. The real reason for the staggering difference is that people only buy Christmas cards before Christmas.

Put more generally, markets contain periodic variation—ranging in granularity from full-blown seasonality to specific weekday or time of day shopping preferences. These periodic forces can

www.dbooks.org

sometimes *fully* account for observed differences in conversion rates between two consecutively run advertising campaigns, as happened with the Christmas card example above. The most reliable way to insulate against such contamination is to run your test variants at the same time as one another, as opposed to consecutively. This is the only way to ensure a fair fight. As far as implementation details go, you can stick your various variants into an A/B testing framework. This will randomly display your different ads, and you simply tally up the results once the experiment ends.

Perhaps you are thinking, "My market isn't affected by seasonality, so none of this applies to me". I strongly doubt that you are immune to seasonality, but for argument's sake let's assume your conviction is correct. In this case, I would still argue that you have a blind spot in that you are underestimating the temporally varying effect of *competition*. There is no way for you to predict whether your competitors will switch on adverts for a massive sale during one week only to turn them off during the next, thereby skewing the hell out of your results. The only way to protect yourself against this (and other) time-dependent contaminants is to run your variants in parallel.

Do not peek at your test results early.

Anyone wedded to statistical rigour ought to think twice before shutting down an experiment after perceiving what appears to be initial promise or looming disaster. Medical researchers, with heartstrings tugged by moral compassion, wish that every cancer sufferer in their trial could receive what's shaping up to be the better cure—notwithstanding that the supposed superiority of this cure has yet to be firmly established with anything approaching statistical significance. But this compassionate rashness can have terrible consequences, as happened in the history of cancer treatment when surgeons performed horrifically painful procedures (such as radical mastectomies) that later turned out to have no benefit whatsoever.

Generally speaking, we should not prematurely act upon the re-

sults of our tests. The earlier stages of an experiment are unstable. During this time, results may drift in and out of statistical significance. For all you know, two more impressions could cause a previous designation of "statistically significant" to be whisked out from under your feet. Moreover, statistical trends can completely switch direction during their run-up to stability. If you peep at the results early instead of waiting until the experiment runs its course, you might leave with the wrong conclusion.

For this reason, it's best practice *not to peek* at an experiment until it has run its course—this being defined in terms of a predetermined number of impressions or a preordained length of time (e.g., after 10,000 impressions or two weeks' time). These finishing posts should be established *before* starting your experiment. If you accidentally happen to view your results before then, resist the urge to act upon what you see, or even to designate them as "facts" in your own mind.

Use a statistical significance calculator.

There are all sorts of technical subtleties that can trip up well-meaning marketers who are foraying into statistical significance. To this end, I recommend the assistance of online tools and programs for calculating statistical significance. To start with, check out Evan's Awesome A/B Tools.[3]

[3]http://www.evanmiller.org/ab-testing/

www.dbooks.org

15 | Email Marketing

Even to this day, email marketing is underrated and underused by online marketers. This may be because we spend so much of our time online that we become desensitized to the content of promotional emails and overwhelmed by their frequency. Empathetically wishing to avoid burdening other people's inboxes, we decide against equipping the businesses of our own creation with email campaigns.

But this is a mistake. Just because you hate receiving promotional email doesn't mean that your customers feel the same way. In fact, promotional emails are welcomed by their recipients when they offer value—be that by announcing the availability of a much-anticipated product or offering a chunky discount for some doodad the customer already planned to buy.

Of course, like anything in life, haters are going to hate. In response to one of our company's emailing campaigns, a disgruntled recipient emailed us to say, "STOP EMAILING ME ASSHOLES". (For the record, I have never emailed anyone anywhere any assholes.) But, for every negative reaction like this, there were at least a hundred sales, completed quietly by smiling members of the silent majority.

Remember too that your subscribers can always opt-out of your campaigns by moving their mouse down to the "unsubscribe" button. Knowing this, you shouldn't be too bashful about sending out emails. Marketing legend has it that the ideal frequency for promotional emails is surprisingly high: about one email per week. That's probably about 50 times more emails per year than most of your competitors are sending.

Email marketing has a compelling commercial case. The initial infrastructure for sending them can be built without serious dif-

ficulty or expense, and the unit cost for each outbound email is practically nil. Once in place, these marketing systems can run with little or no maintenance for years on end, awarding your business a permanent and recurring boost in revenue. It's like having a loyal door-to-door salesmen who works day and night to win you sales.

The strategic goals of email campaigns can be grouped into direct emails that unashamedly ask for conversions then and there and indirect ones that instead lean back and gradually fan up desire. Direct emails have discount codes and links to product pages, whereas indirect emails contain company news, useful tips (which conveniently portray the company well), drip-feed content courses, and notifications about goings-on over on the website. There is much debate about whether direct or indirect email campaigns are better, but little consensus has been reached.

The Law

The sending of promotional emails is regulated by a hodgepodge of region-specific laws such as the CAN-SPAM Act of 2003 (in the US) and Directive 2003/58/EC (in the EU). Generally speaking, full compliance with the main global rules requires that you:

- Ensure that the subscriber opts in. Under EU rules, opt-in consent must be a "(1) freely given (2) specific, and (3) informed indication of the user's wishes".[1] This implies that the ubiquitous practice of automatically subscribing registered users or past customers to mailing lists is illegal. Therefore, marketers who want to play it safe should ask their users to tick an explicit checkbox indicating their consent. Behind the scenes, the play-it-safe marketer should program their server to record the timestamp and IP address for when this ticking happened. This might seem like overkill until you realise that there have been court cases where companies had to provide concrete evidence that an opt-in event occurred. The most risk averse marketers go so far as to use

[1]Directive 2002/58/EC

www.dbooks.org

double opt-ins, whereby potential subscribers must confirm their consent by clicking a link subsequently sent out in an email. Since access to email accounts is password protected, the double opt-in step is seen as strong evidence of genuine consent.
- Provide a valid postal address.
- Provide company information (e.g., registered company name and sales tax number).
- Provide a (working) unsubscribe link within every promotional email.

Even though these requirements are, technically speaking, "the law", it must be said that the overwhelming majority of web marketers ignore all but the need for an unsubscribe link.

Best Practices

- Use images instead of just text for higher conversions rates, but always keep a textual fall back to cater to the substantial number of email users who disable images from displaying in their clients.[2]
- Put a large button in the centre of your email even if it does nothing other than lead to your homepage. This practice is motivated purely by the observation that people love clicking on buttons. The text written on the giant button needn't have any relation to its true purpose. For example, in one of our promotional emails we inserted a huge button that reads, "Activate Discount". The reality is, though, that the discount had already been activated *before* the email was ever sent; we just wanted to give the email recipient a compelling reason to click the button and end up on our website again.
- Dress this giant button up with URL tagging for better Analytics reporting. Ditto for any other links in your emails. Without this precaution, you'll effectively be flying blind vis-à-vis your email performance.
- Use 1x1 pixels (or third-party software) to track open rates and click-through rates.

[2]https://litmus.com/blog/gmail-data-analysis-reveals-image-blocking-affects-43-of-emails

- Pay inordinate attention to your email subject line—unless it's compelling, the email recipient will delete your email without inspecting its contents.
- Be intelligent about who you email so as to avoid bombarding subscribers with irrelevant offers and jolting them to unsubscribe. To be clear, I am not advocating the need for anything remotely as sophisticated as artificial intelligence; in my view, it's sufficient for your campaigns to implement a few simple rules of thumb, such as "people who buy the first book in our series will probably want to buy the other ones" or "people who bought our printer will probably want to order some ink later".

Classic Promotional Emails

Sale

If your profit margins can take it, considering putting on frequent sales. At the very least, sales are a great excuse to remind your subscribers that you exist, and that it's worth dropping by your website again.

Obviously, there's no point in having a sale if no one knows about it. Email marketing comes to your assistance here by helping you spread the word to all your subscribers.

The limited time span of a sale compels leads into buying from you before they lose their chance for a discount. Therefore, be sure to prominently mention when your sale ends. You might even program your email to display the reduced "sales price" for whatever products the subscriber previously browsed.

Cart abandoners

Someone that added one of your products to their cart but didn't finish purchasing it is a grade-A target for continued marketing attention. After all, you know that this person is—or at least *was*—interested.

Enter the cart abandoner's promotional email. The point of this

www.dbooks.org

is twofold: to remind abandoners about the product they liked and to sweeten the deal just enough to persuade them to buy (e.g., by offering a discount). Just as with the sale email, there's no harm in applying a little time pressure too (e.g., "this weekend only").

The trick with discounts in cart abandoner emails is to minimise cannibalising the margins you would have earned had you been able to sell that item at full price. You want to avoid sending discounts to people who would have bought the item in the cart at full price a few days later. One tactic for doing this is to study your historical shopping cart data and figure out the behavioural patterns of your previous cart abandoners. For example, I found in my business that there was a 95% chance that someone would never complete an order 9 days after they abandoned. Based on this data, I programmed my cart abandoners email to only send out a discount 9 days after the cart was abandoned.

Within the email, I present the discount as a store-wide "sale" so as to encourage the customer to buy more that just the original abandoned product, and so as to stop them figuring out that if they leave products sitting in their cart, they will automatically get a discount a few days later.

Recent churners

If your business model is based on monthly subscriptions, you'll know that a certain percentage of your customers end their subscriptions every month. These are good candidates for future marketing emails which remind them to get started again.

16 | Introduction to Paid Advertising

In the following pages, I will cast a bird's-eye view of today's paid advertising ecosystem. What follows is not* a detailed "how to" guide. Instead, I will present an armoury of principles that will steer you clear of the gravest pitfalls in paid advertising and help you navigate the broader landscape. Once you've absorbed these broad ideas, you'll need to search elsewhere for specific implementation details.[1]*

Let's continue by addressing the main objections to paid advertising.

Advertisement Is an Annoyance

Paid advertising is taboo. After an entire childhood of despairing when our cartoons were interrupted by soap detergent advertisements, we've all grown deeply sceptical of anything that smells remotely like paid-for marketing. And for most of us, this reflexive dislike has survived into adulthood.

Effective advertising doesn't feel like advertising to its recipients. It doesn't seem mercantile, pushy, offensive, irrelevant, or even fully recognisable as a sales attempt. Instead, it simply points out that someone out there has a solution to a problem they currently face. Sometimes, those being advertised to don't even *perceive* their having seen an ad. They click a sponsored link and, five seconds later, if you asked them how they arrived on the advertiser's website, they would deny having clicked an ad. It's as if good advertising bypasses people's psychological aversion to commercial messages.

[1]In general, I would recommend watching YouTube videos that were created a maximum of *one year* ago—the video format is easier to absorb and follow when you're learning about where to click within the many menus of labyrinthine advertising platforms.

www.dbooks.org

Imagine you've relocated abroad and need to find a health insurer. You Google for "health insurance" and the top-ranking organic hits point to the websites of big brand insurers. You visit a few of these websites and study their plans and prices. But their offerings are complicated and difficult to compare with one another. This challenge is not made any easier by the sales materials only being available in the tongue of your newly adopted land. You start to tire from the strain of information overload and feel a swell of uncertainty about how to resolve this important decision.

In these situations, what you need is a trustworthy English-language health insurance comparison service. Imagine that one existed, but you never thought to look for it. Its existence wouldn't have done you any good, so you would eventually be compelled to chose a big brand insurer under uncomfortable uncertainty, likely getting ripped off in the process.

Suppose instead that the comparison service had taken out Google advertisements for the search term "health insurance". That would ensure that you would become aware of its services at the time you needed them. Through its website, you would have been able to choose between insurance plans in an informed manner, likely securing a much better deal for this long-term expense.

The effect of the comparison service's advertising was to connect two well-matched parties that would otherwise have never discovered each other. This enriches everyone's lives and thus is the magic of advertising as discovery. The next time a newly opened Indian restaurant (your favourite cuisine) informs you about their nearby branch, think "happy discovery". And the next time the retailer of home cinema projectors informs you that the model you were saving up for is now on sale, think "happy discovery". And when I, the author of this book, advertise my next release, think...

No One Clicks on Paid Adverts

There is a commonly held belief that people searching on Google completely ignore sponsored links and adverts. Any attempt at advertising would remain unseen and ignored—at best, a waste of your time.

How true is this assumption? In 2011, a study by GroupM UK and Nielsen sampled 1.4 billion searches and found that 94% of clicks went to organic results, leaving only 6% to paid advertising.[2] Other studies are less pessimistic: Richard Zwicky from Enquisite studied search traffic patterns for 5000 major websites and found that paid advertising clicks added up to 10.5% of the overall search traffic share.[3]

Although a traffic slice of between 6% and 10.5% is underwhelming, it's worth remembering that a small slice from an obscenely large pie can still feed many mouths. And these little slices are more nourishing than the average because paid search visitors convert at a 1.5x higher rate than organic ones, once they have landed on your website.[4]

There is a demographic trend also worth pointing out: GroupM UK and Nielsen found that while only 5% of those under 17 clicked on paid ads, as many as 23% of those over 55 clicked.[5] The older your target audience, the more likely they are to respond to paid adverts.

To a certain degree, focusing too much on click share is narrow-minded. Even if no one ever *clicks* on your paid adverts, they still *see* them. This visibility establishes your brand in the pub-

[2] https://econsultancy.com/blog/10586-ppc-accounts-for-just-6-of-total-search-clicks-infographic

[3] The original source said, "for every 1 click on a paid search result, the organic results generate 8.5 clicks". I converted this to a percentage for clarity. https://moz.com/blog/via-enquisite-ppc-agencies-make-45x-what-seos-do-for-the-same-value

[4] https://econsultancy.com/blog/10586-ppc-accounts-for-just-6-of-total-search-clicks-infographic/

[5] https://econsultancy.com/blog/10586-ppc-accounts-for-just-6-of-total-search-clicks-infographic/

www.dbooks.org

lic's eyes, which may well lead to future organic business. It's a model not unlike that of old-fashioned newspaper and magazine advertising.

Of course, all the above depends on ad blocker adoption.

Doesn't Everyone Have Ad Blockers, Anyway?

Perhaps the most common objection I hear young, tech-savvy people levy against online advertising is that "most people use ad-blocking software, so paid advertising is dead".

This is a valid concern. To what degree should you be worried?

A 2015 study by PageFair and Adobe estimated that 16% of US internet users are using ad blocking software. Figures for European countries vary from about 10% (France) to 37% (Greece).[6] All in all, ad blockers are still a minority thing.

There are other reasons for advertisers not to feel too disheartened.

Among internet users that have ad blockers installed, most only have them on their primary laptops. Ad blocking for mobile phone browsers is only available in iOS (13.9% market share[7]), and even within this ecosystem adoption is undersubscribed. Furthermore, it's unlikely that Apple or Google will ever allow ad blockers for mobile phone *apps* such as Facebook or Twitter.

Despite what their names suggest, ad blockers do not block all ads. The most downloaded ad-blocking plugin, Adblock Plus, allows unobtrusive and well-behaved advertisements from reputable publishers to pass through its filters. The logic behind their allowing exceptions is that internet users first resorted to ad blockers in response to annoying pop-ups, loud autoplaying video commercials that sucked up bandwidth, and monkey-punch games. Civilised and non-abusive advertisements, like Google Search

[6] http://downloads.pagefair.com/reports/2015_report-the_cost_of_ad_blocking.pdf
[7] http://www.idc.com/prodserv/smartphone-os-market-share.jsp

ads, are far less annoying, so the argument for banning them is less compelling.

Ad-blockage usage is weighted towards certain demographics, and you might be lucky enough to serve an audience with lower adoption rates. PageFair reported in 2014 that while 41% of American internet users aged between 18 and 29 used ad-blocking software, only 15% of those aged over 60 did the same.[8]

Ad blocking also affects some industries more than others. According to PageFair's 2015 report, blocking rates by site category were as follows:

- 26.5% of gaming websites
- 19.1% of social networking websites
- 17.0% of tech/internet websites
- 16.9% of sports websites
- (some left out)
- 5.2% of personals/dating websites
- 5.0% of real estate websites
- 4.9% of charitable organisation websites
- 2.5% of government/legal websites

Paying for Attention Is Dirty

Paid advertising feels wrong for another reason: It seems somehow repugnant to pay for attention, as if buying friendship (or other things).

But this disadvantage is also, in its way, an advantage—particularly for marketers promoting dreadfully boring businesses. In all likelihood, a YouTube video for mortgage refinancing or life insurance won't be going mega-viral anytime soon. Pragmatically speaking, it's wiser for these dull businesses to drop any illusions that their products are interesting and instead just cough up money for attention. Sure, buying a share of the limelight feels fishy, but sometimes that is the price of success.

[8]https://blog.pagefair.com/2014/adblocking-report/

www.dbooks.org

It's Too Expensive

Paid advertising is financially risky: A poorly performing campaign blazes through cash reserves faster than a fire through a city built of nothing but kindling.

New businesses, barring a funding jackpot, are usually strapped for cash. The last thing they want to do is take on more expenses. Naturally, they shy away from paying for advertising...after all, aren't there plenty of free options available, like SEO or social media or press coverage?

Here's the catch: The free options are hard to pull off. The work demanded is labour intensive and, even if you apply yourself and give it your all, the results often lie outside your sphere of control, dependent on capricious search engine algorithms or journalistic opinion that you cannot reliably influence. In general, these alternative marketing channels are free only in the sense that what you save in money, you lose in time.

Moreover, even if your stars align and everything goes supremely well, there can be a long delay before you see any results with a free channel. Paid advertising has the rare and wonderful advantage among online marketing channels of being able to deliver traffic to a new venture *as soon as it opens its doors*. Other channels, such as SEO, can take months or even years before they start delivering respectable amounts of traffic. Many new businesses cannot afford that kind of lag. I know I couldn't when I started out. Today, my business Oxbridge Notes receives well over 90% of its traffic from SEO, but when I launched, the figure was stuck firmly at 0%. Back then, Google Adwords sent me all my traffic, and I have them to thank for my first few hundred sales. In fact, my business wouldn't have survived long enough for me to be writing this book if it weren't for paid advertising shuttling me through those first few patchy months. Although SEO is ultimately more efficient in the late game thanks to its low per unit cost, paid advertising is a necessity at kick-off.

Compared to other approaches, paid advertising is (relatively) ef-

fortless, fast-acting, and tractable to your will. It's not unrealistic to have your first campaign up and running within a few hours, and to receive your first sale that same day.

And as regards worries about financial costs, there is one simple formula to remember: As long as your unit advertising costs don't exceed your unit profits, paid advertising remains a good deal. And it becomes an exceptionally good deal if your unit advertising costs are only a mere fraction of your unit profit.

Why Pay for Traffic That Would Have Visited Anyway?

It's a common concern amongst internet advertisers that they might pay for traffic they would have otherwise gotten for free through other channels.

Imagine a tattoo artist with an established and effective SEO campaign. Anyone in their city who googles "tattoo artist" sees the artist's website rank #1 in the organic results. Given the artist's existing success in organic search, is it worth their while advertising on the same keywords with Google Adwords? Isn't there the danger that whoever clicks on the tattoo artist's paid adverts would have clicked on their free organic results if there were no paid advertisements jostling for their attention?

There is some truth to this fear. Without a doubt, a percentage of the tattoo artist's otherwise free traffic will end up as paid-for traffic, as least when the artist advertises on keywords which are already receiving organic traffic. That said, advertising on promising *search terms for which the artist barely ranks* is a good idea though. These terms are essentially brand new territory, and the artist won't be competing with their organic results by advertising on them.

Despite the toll levied by paid advertising, I would still argue that the artist should go ahead and pay for advertisement. My grounds are as follows:

www.dbooks.org

1. The cost isn't that high.

At least with Google Adwords', their Quality Score system dramatically reduces the price of advertising when there is, in their eyes, a good "match" between the keyword targeting and the page the advertisement leads to. In practice, I have often noticed a cost per click of £0.01 when advertising for any keyword I already rank #1 for in organic search.

2. Countermeasure to competitor advertising.

At any time, it's possible for competitors to start advertising on keywords you rank #1 for in organic searches. Their campaigns poop on your party and siphon away leads before they get as far as reading the organic results where you shine. When your competitors start advertising, often you need to respond with advertising of your own just to retain your market share. This is especially true when your competitors avail themselves of special paid-ads-only bells and whistles. For example, Google Adwords has a feature called "ad extensions", which are extra widgets within the advert that draw the eye and encourage conversions. We'll discuss these in a later chapter.

3. Control over advertising messages.

In organic search, Google decides what text snippet accompanies your result entry. They are *influenced* by an on-page SEO element known as the meta-description, but not bound to obey it. All too often, Google completely ignores this meta-description directive and just displays a random jumble of words taken from your page. By way of contrast, paid advertisements give the marketer full control over the exact wording of their adverts, enabling them to craft their text to present themselves exactly as they please.

And *even if* Google Search completely respected the meta-description directive, Google Adwords would still offer superior control because advertisers can create different advertising copy variants for every single keyword advertised upon. For example, the keyword "henna ink temporary" should be matched with an advertisement whose copy highlights how long henna ink lasts,

whereas the keyword “tattoo pain” would be better focused on options for minimising pain during the tattoo session. With this kind of flexibility, the tattoo artist can directly address the concerns of their leads as indicated by what they typed into Google. The net result of this adaptivity will be higher click-through rates.

Organic marketers cannot avail themselves of this same flexibility because each webpage is limited to one meta-description. Every web surfer, no matter what they searched for in Google, will be presented with the same unchanging and hopelessly general front (that is, unless Google decides to disregard the meta-description for that search and instead throw up gobbledygook).

4. Social signalling of trust/credibility.

Business-minded customers believe that when a company buys paid advertising, that counts as a rough indicator that the company is an established and profitable entity, and one that is therefore reliable and safe to deal with. The act of buying advertising is a type of social signalling that establishes seriousness and credibility.

(Side-note: In some markets, adverts may have the opposite effect. Anti-“The System” types may never click on adverts out of principle. They might also believe that any company paying for a public boost must be compensating for the substandard quality of their product.)

www.dbooks.org

17 | Paid Advertising: First Steps

Building an "Advertising-Ready" Web Application

1. Create appropriate landing page constellations

Many new advertisers make the mistake of sending all their paid traffic to generic homepages. This is a bad idea because it is unreasonable to expect rushed leads to navigate your website until they see the product they want. Most internet users are impatient, so this strategy is really just a recipe for wasted advertising money and heavy use of the back button.

It's a much better idea to send leads directly to pages corresponding to the advertisements they were shown. Usually this means directing traffic to the appropriate product pages—but not always.

For one, having different landings pages for each product won't make any sense if the business is unable to tell, upfront, which *exact* product the customer wants. Instead, the business should create landing pages for each *category*. These pages will then help the customer to choose the right product for them. For example, Amazon might advertise a page showcasing their various in-ear headphones instead of arbitrarily bringing traffic to some particular headphone set or other.

Other businesses, such as those that offer only a *single* product or service, should build separate landing pages that each address different subsets of customers. For instance, a website selling translation services might have separate landing pages for each of the following: "translation of legal contracts", "translation of instruction manuals", and "translation of movie scripts". These specialised landing pages would focus on persuading specific sets of customers to convert, without having to be everything to everybody.

For all this segmentation to work, the website has to have been built with separate landing pages for the different things that customers might want. And it's in this sense that the website needs to be "advertising-ready".

Creating all these landing pages has a double pay-off: As well as sharpening the effectiveness of paid advertising, keeping a lively cast of landing pages also boosts SEO. The numbers here are quite astonishing: Hubspot, after studying 7000 websites, found that companies with 15 landing pages generated 55% more leads than companies with only 10.[1]

2. Calculate profitability with code

As we'll see later, it's important to know your unit profitability when advertising. The math of profitability would be cumbersome and error-prone to calculate by hand or calculator, so instead it should be calculated automatically based on information stored in the database, such as inventory costs, postage costs, sales tax costs, financial fee costs, and royalty costs. By writing code that calculates unit profits on a per-product basis, businesses selling many different products will be armed with the knowledge necessary to throttle advertising budgets for each item and thereby optimise overall profitability.

3. Write automated tests for conversion pixels (etc.)

Modern online marketers advertise with not one advertising platform, but many. As you add additional platforms, you end up with more and more conversion pixels installed on your website, and this increases the chances that one of them will break and start under- or over-recording conversions.

This is a particularly damaging case of incorrect data because it will cause you to make poor business and spending decisions that will hamper profitability.

In regards to the specifics of how to test conversion pixels, that's

[1] http://blog.hubspot.com/blog/tabid/6307/bid/33756/Why-You-Yes-You-Need-to-Create-More-Landing-Pages.aspx

www.dbooks.org

for the programmers to decide. (My ghetto approach is to use integration tests to scan for the pixel presence at the key moments and then supplement these automated tests with manual tests after large changes.)

4. Only advertise to permanent URLs

Online advertisements deliver traffic to specific pages of your website. If that page disappears, the advertisements stop working. And trust me, there's nothing more annoying that investing a day into writing a multi-platform set of online advertisements, only to discover that they've broken down a few weeks later because the URL they were supposed to lead to no longer exists.

URLs change for all sorts of reasons (e.g., refactors by your programming team or content edits by users of your website). To avoid hassle, ensure that you exclusively advertise to URLs that aren't subject to future changes. Personally, I create separate unchanging "permalink" URLs for any content I want to advertise to, shielding them from breaking edits.

5. Place conversion pixels high up in the HTML of the page

Just because there is a conversion pixel on a page doesn't guarantee that it will always register conversions taking place. Amongst other things, it's necessary that the browser execute the code requesting the pixel. Because browsers execute HTML code from the top of the file to the bottom, conversion pixels low down in the file may never be fired. This could happen, for instance, when the customer quickly skips to the next page after converting.

Analytics and Advertising: The Eyes to See With

Before the computer revolution, businesses—which back then advertised with newspaper, TV, or other classic media—could only gauge their marketing success with the roughest and most rudimentary of tools. Mostly, they switched on adverts for one month and calculated how the subsequent sales compared to previous periods where advertising was switched off. Measuring success was an almost mystical undertaking, a task caked in deep,

unsettling uncertainty.

Imagine a pet food seller. It opens its doors in December and starts taking its first orders. Without advertising, it achieves the following results:

Pet Food	December Sales
Dog Food	£3,000
Cat Food	£4,000
Total	£7,000

In January, the pet food seller buys advertising in two different local magazines—one for dog lovers and another for cat lovers. It expects to drastically increase sales in January, but at the end of the month its accountant tallies the results and they turn out to be underwhelming.

Side-note: *I use the standard accounting convention of printing negative figures within parentheses.*

Table 17.2 – "Increase" is the additional sales made in Jan relative to Dec. "IOAS" stands for "Increase on Advertising Spend" and is calculated here as "Increase in Sales - Advertising Costs". It is only a coincidence that both pet food categories had the same amount of sales in Jan.

Pet Food	Ad Costs	Jan Sales	Increase	IOAS
Dog Food	£2,300	£6,000	£3,000	£700
Cat Food	£3,000	£6,000	£2,000	(£1,000)
Total	£5,300	£12,000	£5,000	(£300)

At first glance, the data appears damning. Although the advertisements did increase *sales*, this increase came at the cost of wiping out *profits*, mostly due to the abysmal performance of the cat food advertising campaign. It seems like advertising was a really bad

www.dbooks.org

idea.

But there is a problem with our analysis: Our calculations assumed that the pet food seller's December sales figures were the *baseline*, and we judged the effectiveness of January's foray into advertising relative to this assumed base rate. This assumption may not be justified. What if December was an outlier month, one with *far more* sales than could be expected in other months? Perhaps it's a seasonal trend whereby pet owners buy overpriced Christmas treats for their dogs and cats. Or maybe it's a press coverage thing, with the December boost in sales attributable to a local magazine doing a write-up about the pet food seller's launch.

Let's assume that we had access to an alternative universe where the pet seller had forgone advertising in January. What would it have earned?

Table 17.3 – "Jan Counterfactual" is the key column here, representing the amount of sales that *would have* happened in January if the pet seller took out no advertisements. I copied over some figures from the last table (i.e., the world *with advertising*) so you don't have to flick back and forth.

Pet Food	Dec Sales	Jan Ad Costs	Jan Sales	**Jan Counterfactual**
Dog Food	£3,000	£2,300	£6,000	**£1,500**
Cat Food	£4,000	£3,000	£6,000	**£2,000**
Total	£7,000	£5,300	£12,000	**£3,500**

With the benefit of this counterfactual analysis, we now see that the return on £5,300 of advertising spend would be substantial: It would be responsible for an £8,500 boost in revenue (= £12,000 - £3,500), representing a return on advertising spend of £3,200 (= £8,500 - £5,300). Had we been able to peep at this alternate universe where the pet food seller hadn't advertised in January, we would have known that the advertising was actually plenty profitable.

As the above example has shown, it was often impossible for busi-

nesses advertising in traditional media to say with any confidence whether an advertisement was profitable (at least without a time machine). Ultimately, this problem stemmed from their inability to connect and attribute individual sales to acts of advertising. But with the advent of online advertising and internet technologies, this changed. There appeared all manner of sophisticated new tools for gauging advertisement success, the most important of which was conversion tracking.

Conversion tracking

Conversion tracking is the technology that obviates the need for a time machine by providing a traceable link between sale and advertisement. Whenever a lead sent by an advertising provider buys something, conversion tracking attributes the sale to the specific advertisement (and advertising provider) that was ultimately responsible.[2] These technologies, when properly implemented into a business's practices, end the guesswork about whether advertising was worthwhile. The pet food store would no longer wonder how many of its January sales would have happened "naturally" without advertising; instead, it could simply read these figures off an automatically generated chart on the advertising platform's website.

Table 17.4 – "Ad Orders" refers to orders that were attributable to the action of advertising.

Pet Food	Ad Cost	Jan Sales	Natural Orders	Ad Orders	Total Orders
Dog Food	£2,300	£6,000	95	5	100
Cat Food	£3,000	£6,000	5	95	100
Total	£5,300	£12,000	100	100	200

From the above table, we see that the overwhelming majority (95%) of cat food orders were attributable to advertising, whereas

[2]Some marketers have different conversion goals, such as mailing list sign-ups, app installs, or whatever else the website owner considers winning at business. For the sake of simplicity, I'm restricting the present analysis to sales.

www.dbooks.org

practically none of the dog food orders were. It appears that the cat food advertising campaign was way more effective. Is it time to stop the analysis and switch off the dog food campaign?

No, it isn't. The number of orders is largely irrelevant; what matters is sales (or revenue). Judging an advertising campaign's effectiveness by the number of orders it triggered assumes that every order was for an equal amount of money. This won't be true if, for example, the advertising campaign managed to reach professional buyers who bought in bulk, whereas the orders coming "naturally" were placed by individual consumers who only bought tiny amounts. Our tentative earlier conclusion to switch off the dog food advertising campaign because it didn't generate enough *orders* would be wrong if those few dog food orders nevertheless accounted for a disproportionate sum of the total *revenue*.

The antidote to this uncertainty is to recognise that not all conversions are created equal. This is done by assigning specific financial values to each order. Instead of just recording the number of orders received, the pet food seller would now record the total revenue generated by these orders.[3]

Table 17.5 – "Avg Ad Order" refers to the average basket value of orders placed by customers found through advertising. To save you from the need to refer to the previous table, I'll repeat that there were 5 Ad Orders for dog food and 95 for cat food. "Ad Revenue" refers to the revenue attributable to these same customers. It is equal to the Avg Ad Order times the number of orders placed by these people. "Ad Profit" is calculated by subtracting the Advertising Cost from the Ad Revenue.

Pet Food	Ad Cost	Jan Sales	Avg Ad Order	Ad Revenue	Ad Profit
Dog Food	£2,300	£6,000	£800 (wow!)	£4,000	£1,700
Cat Food	£3,000	£6,000	£47-ish	£4,500	£1,500
Total	£5,300	£12,000	£85	£8,500	£3,200

[3]Google Analytics has a special plug-in designed for doing just this: The Ecommerce Plug-in https://developers.google.com/analytics/devguides/collection/analyticsjs/ecommerce

This improved data invalidates our previous tentative conclusion that the dog food advertising campaign was wasteful. Because of the large difference in order size between the groups, the dog food campaign was worth running, and its results turned out to be slightly *better* that those of the cat food campaign, despite its many individual orders. Talk about a comeback.

Now that we know from our conversion tracking that *both* the pet food seller's advertising campaigns generated revenue that would otherwise not have been earned, should it feel confident to continue advertising?

Just when you thought I had exhausted my repertoire of uncertainties and turnarounds, I have one more to introduce: bottom-line profit. When we concluded that both advertising campaigns were worth keeping, we assumed that all revenue was worth the same amount to the business. But this isn't true—the figure that *really* matters is bottom-line profit.

The pet food seller has variable costs (costs that change in proportion to the amount of pet food sold), such as the price of raw ingredients, packaging costs, shipping costs, and sales tax. For our purposes, let's define profit as the difference between the price the food sells for and its variable costs. That gives us:

Pet Food	Sales Price	Variable Costs	Profit
Dog Food	£20	£10	£10
Cat Food	£20	£3	£17

The above table shows us that £20 of revenue from cat food sales is worth more in profit than the same amount of revenue from dog food sales. Because the pet food seller is ultimately optimising for profit, it would prefer to sell more of its high-profit item—the cat food—than the less-profitable dog food. Therefore, it should pass profit figures to its conversion tracking rather than just revenue numbers. If it can pull this tracking feat off, the pet

www.dbooks.org

food seller will have nicely informative data for making profitable business decisions:

Table 17.7 – Shows the cost of the extra goods sold thanks to advertising.

Pet Food	Ad Revenue	Cost of Goods
Dog Food	£4,000	£2,000
Cat Food	£4,500	£675

Table 17.8 – "Ad Profit" again refers to what's left after subtracting the advertising costs from the revenue generated thanks to advertising. "Bottom-Line Profit" then subtracts the Cost of Goods from this figure.

Pet Food	Ad Cost	Ad Profit	Cost of Goods	Bottom-Line Profit
Dog Food	£2,300	£1,700	£2,000	(£300)
Cat Food	£3,000	£1,500	£675	£825
Total	£5,300	£3,200	£2,675	£525

Now we have our final turn-around: Everything considered, dog food is actually *losing money* through advertising, whereas cat food is the campaign that is, ahem, bringing home the bacon.

When conversion tracking statistics still err

Because of the limits of modern tracking technologies, there will always be some inaccuracies with the recorded business intelligence figures.

In order for us to gauge the success of an advertisement at causing sales, there must be an unbroken trackable link between the initial advertisement and the ultimate conversion (e.g., the sale). But this link can get severed, for example, when someone first clicks your advertisement on their mobile device and decides not to buy right then, but later googles you on their laptop and completes the sale. Even with the latest technological precau-

tions,[4] these two visits may appear separate and unrelated within your conversion tracking reporting, which will cause the effectiveness of your advertisements to be underreported. (The conversion here would be attributed to organic traffic.) Upon viewing this lacklustre performance in your reports, you might be provoked into ditching your advertising campaigns. But doing so would trigger a greater decrease in sales than your conversion data would lead you to expect because you will lose out on all the sales that your tracking was unable to properly attribute.

There are myriad other ways that the attribution link might get severed. Imagine a touring musician advertising their upcoming concert to you. Upon seeing their ad, you decide against buying tickets online, and for buying them in person at the doors. With no electronic trail connecting the online advertisement to your eventual ticket purchase at the venue, the ad will appear a flop, even though the ticket sale would never have happened without it.

The fruits of advertising sometimes manifest long after the ads have finished running. An ad that causes no sales today may nevertheless leave a subtle impression on whoever sees it. This past exposure leads to future affinity—months later, when deliberating over which product to buy, this same person may feel inexplicably inclined towards the product they once saw advertised. These staggered benefits will not appear in recent conversion statistics.

The offshoot of all the above complications is that we usually underestimate the benefit of paid advertising. As such, while we should pay attention to our conversion statistics, we should not be slaves to them. Despite all the technological advances, we should still keep a watchful eye on how overall sales figures change as we switch campaigns on or off. The technology has improved since the days of newspaper ads, but it is still far from perfect.

[4] Cross-device analytics wizardry, as described in the Google Analytics chapter, helps lessen this problem.

www.dbooks.org

The guiding rule: Know your margins before advertising

As we saw in the previous section, performance data can misrepresent the truth when it is shallowly analysed. The only way for a business to be on solid ground is to track *profit* figures and base their advertising analysis and spending off of these.

If a single bag of cat food has a pre-advertising profit of £17 for a £20 bag, that means the pet food seller could pay up to £16.99 on advertising each bag and still remain profitable.

In practice, of course, the pet food seller won't pay £16.99 for advertising each bag because they would only be left with a mere penny profit on each unit sold. Instead, the seller would expect to earn a minimum margin for itself on each bag (e.g., £2.00). In this case, the most the pet food seller would pay in advertising would be £15.00 per bag sold. Obviously, it'd be better for the seller to pay less in advertising costs because that would leave more profit behind; £15.00 is just its absolute upper limit, constrained by the seller's internal costs and the minimum profit it is comfortable keeping for itself.

The business that knows its own costs and desired profit margins has an incredibly simple strategy to guide all their advertising decisions: minimise the amount paid per conversion and never continue any advertising campaign that doesn't leave a comfortable profit margin. This is the most important guiding rule to remember.

Pay Per Click? Per Impression? Per Conversion?

Online advertising platforms offer various different bases on which they will charge you for their advertising services. Are you better off when the payment is reckoned in terms of impressions, clicks, "actions", target ROAS ("Return On Advertising Spend"), or installs?

Before answering this question I would like to draw a distinction between a) the basis on which you *evaluate* advertising success;

and b) the basis on which you *bid* for advertising.

The basis on which you *evaluate* advertising should always be the one closest to your business goal of profiting. Thus, when possible, you should use ROAS ("Return On Advertising Spend"), i.e. the amount of profit you made on your advertising spend. (When this advanced metric is not available, I would evaluate based on cost per conversion.) The reason for choosing these as evaluation bases is that they correspond closely with how we think about our businesses and therefore our analyses become easier to carry out. Case in point: without knowing anything about my business, which of these sounds like it's a better deal: "£0.40 per thousand impressions" OR "200% return on advertising spend". Obviously, the return on advertising spend sounds like the safer bet–it's easier to understand. If we wanted to evaluate whether the cost for thousand impressions was a good deal, we'd have to do a ton of math about the likelihood of each impression clicking and then converting. Yuck.

That said, as much as I prefer *evaluating* on these high-level bases, the question of what basis I *bid* on is determined by what is cheaper on the advertising platform. I found that advertising costs can sometimes be dramatically reduced by asking the advertising platform to charge me on a different basis: for instance, I remember a Facebook campaign three years ago which cost me about £0.80 per click when I bid on a "per click" basis. Behind the scenes, my evaluative basis was "cost per conversion" and this cost per click worked out at about £30.00 per conversion. When I switched to bidding "per thousand impressions", my effective cost per conversion, calculated by looking at the number of conversions the advertising generated, dropped to about £14.00. I don't pretend to know the reasons why the choice of bidding basis makes a difference to advertising costs; I simply observed that it does.

The above example isn't intended to show you that bidding by impressions is cheaper; sometimes it is and sometimes it isn't. Rather, my point is that the basis for *bidding* is worthy of your ex-

www.dbooks.org

perimentation because different bases may lead to reduced costs when compared within a common *evaluative* basis.

18 | Paid Advertising: Staying Organised

It's altogether too easy to become overwhelmed with your online advertising efforts. The accounts you create on various advertising platforms gradually fall into a state of disorder and disrepair. Without active measures to stay on top of things, you'll one day reach a point where you'll dread modifying even the tiniest aspect. For me, it once got so bad that I would procrastinate maintaining my advertising accounts by instead filling out my tax return, which I considered at the time to be the lesser of the two evils. Predictably, this period of sustained neglect caused my advertising performance to dip into free fall.

Structuring and Labelling Your Accounts Sensibly

Online advertising platforms come equipped with somewhat standardised tools and terminology for structuring your accounts. Typically you have a couple of "campaigns", which contain a few "ad groups", each of which is itself composed of a selection of different "ads" and is "targeted" to a certain audience. Let's continue by connecting the abstract structure just given to how it gets used in practice.

Campaigns

Starting from the highest echelon of this organisation chart, campaigns most commonly correspond to *business goals*.

For instance, an online retailer that sells shoes and socks would normally have one campaign for each major category of products: one for selling shoes and another one for selling socks. If the retailer starts stocking(!) some other class of products (e.g., shoe polish), then it just creates a third advertising campaign. Now, if

www.dbooks.org

the advertising for one product class tanks, the retailer can switch its budget to a more successful campaign.

The business goals of a two-sided marketplace (as expressed through advertising) normally reflect the priority such businesses give to recruiting people for both the supply and demand sides of their markets. For example, a startup building a marketplace for piano lessons would have one campaign to recruit piano teachers and another to recruit piano students seeking lessons. This way, advertising budget can be shifted towards the side of the market that most needs a boost.

Other businesses divide their efforts between branding and getting sales. Thus, a luxury watch manufacturer would have one campaign corresponding to gaining admirers and another to sell watches at their store.

The names of advertising campaigns should reflect the business goals behind their creation; for example, good campaign names for the above examples would be “Sell Shoes”, “Sell Socks”, “Recruit Piano Teachers”, “Recruit Piano Students”, “Build Brand Awareness”, or “Sell Watches”.

One advertising platform (e.g., Facebook or Google Adwords) is rarely enough these days. Instead, you will likely create parallel campaigns for the same goal with various other advertising providers. Make sure you keep the campaign names the same across all these platforms. If the campaign for recruiting piano teachers is called the same thing on Twitter, Facebook, and Adwords, it is much easier for you to turn it on and off, to change its overall budget, or to compare how well each platform performs at reaching it. Without this kind of parallelism in naming, managing campaigns spread out across many platforms gets thorny and becomes unacceptably error-prone.

Ad groups

Moving down to the middle-management level of account structure, we find ad groups.[1] These typically correspond to two ideas that often show up inseparably intertwined: 1) targeting options and 2) finer-grained versions of the business goal of the ad group's holding campaign.

With Google Adwords, advertisers usually create one ad group for each cluster of related keywords, such that different ad groups are triggered depending on the keyword someone searched for. For instance, a shoe retailer would have ad groups corresponding to each individual product (e.g., UGG boots or Birkenstocks).

Of course, some businesses only have a single product/service on offer, as would be the case with a furniture transport company. Its ad groups should instead correspond to specific customer concerns as ascertained through the keywords their customers typed. Google queries about "insured furniture transport" ought to trigger ad groups highlighting the company's comprehensive insurance coverage, whereas queries about "last minute furniture transport" should trigger ad groups that emphasise the transport company's quick response time and punctuality.

Due to differences in available targeting mechanisms across platforms, Facebook's ad groups (or "ad sets", as they call them) may instead be associated with targeting options other than keywords, such as the social network's members' relationship statuses or the subjects they studied at college. As such, the retailer selling university textbooks might create an ad group corresponding to each individual field of study—one for math, another for law, another for medicine, etc. Through using the available targeting options, the retailer breaks down their campaign goal into focused subunits encapsulated in ad groups.

With this connection between ad groups and targeting in mind,

[1] Ad groups have been around for a while in Facebook and Google Adwords, but are relatively new in Twitter. At the time of writing, you couldn't modify Ad Groups in the Twitter UI, and instead had to modify them through exported and imported spreadsheets. I would guess that Twitter will update their UI by the time you read this though....

www.dbooks.org

I believe the most helpful way to name ad groups is according to their targeting. For instance, a language translation service would have ad groups named "Legal Translation (T:keywords)", "Movie Script Translation (T:keywords)", and "Interpreter Service (T:keywords)".

Notice how I included *the mechanism* of targeting within the ad group names, communicated with "T" (for "targeting") and followed by the specific targeting mode ("keywords" in this instance). The extra targeting information is helpful to have in the name because otherwise we'd be unable to tell the difference between ad groups that advertised the same thing via different targeting options. Suppose we advertised law textbooks through three different targeting mechanisms: one that targeted people who studied law at university, another that targeted people who interned at law firms, and another that targeted people who wrote about law on their timelines. Without differentiating between these targeting options within our ad group names, our performance reports wouldn't communicate how effective each targeting option was at its advertising job. But this is a critically important data point in marketing intelligence. The choice of targeting mechanism is a huge determiner of advertising effectiveness. Perhaps targeting law *students* was a waste of money, but targeting law *interns* was pure gold. Without the suggested naming convention, the only way we could answer these basic questions of effectiveness would be to open up the settings of every ad group on analysis day and painstakingly compare targeting options so as to determine the differences. A little upfront organisation would have saved us a whole lot of chaos.

As with campaigns, it's important to work on keeping inter-platform and intra-platform consistency in naming ad groups. There's not much to say here, other than advising advertisers not to *guess* what they called an ad group in another platform, but rather to go back and look it up.

In this section I showed you a slightly idealised scenario. The truth is a little messier, mostly because Google Adwords fails to

provide a clean separation between campaigns and ad groups along targeting lines. The problem is that *most* of the targeting options in Google Adwords are set at the ad group level (consistent with other online advertising platforms), but *some* targeting options are set only at campaign level—for example, geographic region. If an advertiser on Google Adwords wishes to show different ad groups to people in different locations (e.g., "UGG Boots" only to USA residents and "Birkenstock" only to German and French residents), then they would have to have two separate campaigns. I would advise the advertiser to name these campaigns so as to communicate the location within the name (e.g., "Recruit Piano Teachers L:USA" and "Recruit Piano Teachers L:Germany&France"). ("L" denotes location.)

Ads

At the lowest level of the campaign structure are the individual adverts. These contain the actual creatives that the public sees. Every ad group must thus contain at least one ad, otherwise there's nothing to display when the ad groups get triggered. But I wouldn't recommend putting *more than* one ad in each ad group, unless you are consciously serving one of these two purposes: 1) protecting against ad-blindness by refreshing your creatives, or 2) optimising the creatives themselves (as opposed to optimising your targeting, which is done at the ad group level, and which generally gives you more bang for your buck).

Not all platforms let you name your individual ads, but on the ones that do, use the opportunity to communicate how the ads are unique with respect to their headlines, bodies, or images. I have a naming convention whereby "I" stands for image, "H" for headline, and "B" for body. By following it, I named one ad "I:coffee-face,H:better-grades,B:instant-download", another "I:clever-cat,H:better-grades,B:instant-download", and another again

"I:coffee-face,H:better-grades,B:free-samples". This naming convention helped me notice patterns about what images, headlines, and copy bodies beat the averages, even when I had hundreds of

www.dbooks.org

combinations to compare.

Don't Delete Data

Once a marketing campaign has run its course, there's a temptation to delete it on the platform hosting it so as to declutter your account. Resist this temptation. Whenever you delete advertising data, you also delete the original records and results of past advertising forays. You might believe that this data contains no further insights, but this would be a premature and somewhat self-centred conviction.

The conviction is premature because future improvements in your skill as an advertiser (or future enhancements to the reporting capabilities of advertising platforms) may enable you to learn new lessons from old data—assuming that old data is still around. Furthermore, deletion of data is self-centred because someone else without your direct experience may one day inherit your advertising account. Without your old data to go on, your replacement may be doomed to repeat many of your past mistakes.

Many advertising platforms let you hide deleted campaigns from the UI, which should be enough to declutter your workspace without the need to permanently remove the data. Where this hiding feature isn't available, there may be the possibility to export your old data for future reference. With a backup handy, there is no issue at all with deletion.

Change History

When the engineering team modifies the website source code, they rely on specialised "source control" systems (e.g., Git or Subversion) to preserve a precise record of how the code used to look and what their changes were. (These source control systems are like the revision history features of Wikipedia or Google Docs, except pumped up on state-of-the-art feature steroids.) Thanks to source control technology, your engineering team can surgically roll back problematic changes whenever a bug is detected.

Instead of struggling to remember how the previous code looked, they see a side-by-side comparison on their screen.

Sophisticated online advertising campaigns can reach complexity levels that rival or even exceed the source code of the websites they promote. The slightest change to a campaign can upset a delicate advertising balance in ways that only become apparent weeks after the incident. As such, there is a parallel need for source control technologies in online advertising. This, luckily, is starting to be provided by some of the big platforms. Google Adwords has an account history feature that lists all modifications made by any team member (complete with 30-day undo),[2] and Facebook has something similar (minus the undo button).[3]

When these features are not available on an advertising platform (or when they are but you want to supplement them to be thorough), I would advise you to create your own change log for jotting down the most momentous of changes. Within this log, you would note down not only the changes made, but also the reasoning behind those changes. Without some means of retracing your steps—both physical and mental—there's a risk that your advertising campaign will one day cease to be effective and you won't have a clue how to repair it. The log will then be to your advertising campaign what the trail of breadcrumbs was to poor Hansel and Gretel.

Schedule Review in Future

Always, always, always schedule a future review for checking up on your online advertisings changes. As you close the tab for today's session with your online advertising platform, pull out your calendar right away and book a review date with yourself a week or so from now. This isn't negotiable—the changes you just made to your campaign will affect it in ways that you can't possibly anticipate, no matter how good you think you are. Without reviewing the changes, you will wrongly presume that everything

[2]https://support.google.com/adwords/answer/2454137?hl=en
[3]https://www.facebook.com/business/help/289211751238030

www.dbooks.org

worked out as expected—even when it didn't.

For instance, I've made "improvements" to my Adwords campaigns, presuming they were for the better, only to log in months later to find that I'd chewed through thousands of euros in waste. Don't be like me.

19 | Paid Advertising: Targeting

Perhaps the most basic goal of any advertisement strategy is informing potential customers that you exist. But not every Tom, Dick, and Harry out there is a potential customer for your particular business, so it's important to direct your advertising towards only the most relevant and promising of leads. This is the art of targeting. To wit: The solar panel seller who takes out advertising in a renewable energy monthly magazine is sure to win far more sales than the competitor who spends the same advertising budget driving around housing estates and announcing the solar panels with a megaphone.

Good targeting saves you money because you don't waste your advertising budget on people who have no interest in what you are selling. Good targeting respects the public's time because they will be subjected to fewer messages about irrelevant products and services. Good targeting treats the public more humanely because you are less likely to bombard the infertile and the childless with advertisements for nappies and children's swimming lessons. Good targeting even helps the platforms you advertise with (such as Google AdWords) because their users are more likely to click on ads and thereby generate revenue for the platform. And because advertising platforms appreciate when you make them money, they reward you by further reducing advertising costs. (See, for example, Google AdWord's Quality Score).

I believe your choice of targeting mechanism is the most important contributor to your online advertising success, more so than the advert creative. Finding the right targeting combination is a bit like panning for gold—the best advice is to try out all sorts of weird and wacky ideas until you stumble upon a keeper.

The modern interface of an online advertising platform is a fighter jet cockpit filled with a bewildering array of controls for

www.dbooks.org

honing in on your potential customers. These controls are mostly filters that limit your advertising to only show in certain geographical locations at certain times of the day, to people who fall into certain demographics or exhibit certain behaviours detectable by the advertising platform.

Knowing how to zone in on your ideal customers from the cockpit of one platform is no guarantee that you'll be able to locate them when piloting another. Each platform is equipped with differing controls, so you may need to test drive a few different models until you find a platform that adequately meets your needs.

Indeed, there is already some implicit targeting going on as soon as you choose to advertise with platform X rather than platform Y. Already, the universe of people who might see your adverts has been narrowed down to "people who use the internet AND happen to congregate on Platform X". If the demographics of Platform X are less than well matched with those of your product's market, this could be quite the unwelcome restriction.

The most fundamental thing that anyone studying targeting ought to learn is the inventory of targeting options available on each platform. For instance, suppose you have a business selling textbooks to medical students. If you know that Facebook lets you target by university degree, then your work is mostly done and you will do admirably well. If, on the other hand, because of inexperience with Facebook's platform, you are unaware of this targeting option, then you will struggle with less-effective targeting options and flounder at the heels of better-informed competitors.

Clearly, it pays to stay aware of what targeting options are on offer. For that reason, the second half of this chapter is devoted to giving you a high-level guided tour of what's out there right now. But this guide is neither comprehensive nor guaranteed to stay up-to-date. You would do well to supplement it with an intuitive rule of thumb that applies whenever an advertising platform is built on the back of a regular website/social network (e.g., Facebook/Twitter): Expect that any data gathered from the website's

users will eventually reappear as a targeting option for advertisers. For instance, Facebook asks its members about their relationship statuses. This data is then fed to advertisers in the form of a targeting option.

The gap between your *intended* audience (the people out there who would love your product and to whom you want to advertise) and your *targeted* audience (the people who actually see your advert as a consequence of your chosen targeting options) is a constant source of frustration and disappointment for advertisers. Ideally, you want both groups to be one and the same so that you advertise to all the right people—and not another soul. But that's not going to happen. There will always be waste—sometimes more and sometimes less. Instead, we must live with calculated levels of imperfection. In some cases, it'll be worth advertising to a group where only 1 person in 100 is in your *intended* audience. This will be so when you reckon that the costs (financial or otherwise) of advertising to this group are lower than the expected profits. The logic is the same as that of the blackjack player who counts her cards and places a bet whenever the odds are in her favour.

The brain of the marketer thinking about targeting is filled with Venn diagrams. Every additional targeting option that they apply adds another circle into the mix. The marketer then adjusts and moves these circles until the shaded area paints out a space as dense with potential customers and as sparse with waste as can be achieved.

General Principles

Off-target targeting

The controls of your cockpit are hardly those of a Rolls-Royce; the levers are rusty, the gears sometimes jam, and the numbers on the speedometer are off by a worrying amount.

There is no guarantee that your customers are as they present themselves on Facebook or Twitter or any other website. And

www.dbooks.org

because most advertising platforms base their targeting on data their users input, we advertisers ought to stay vigilant about the difference between how our customers are in actuality and how they appear within the house of mirrors that is the advertising platform.

The sources of inaccuracy are many: privacy concerns, public posturing, laziness, or staleness (e.g., someone's circumstances change but they don't update Facebook). Facebook has this author on file as a 90-year-old man who is a fan of trashy reality TV shows, but only one of these two "facts" is true.

I don't want to exaggerate the problem of dodgy data because a sizeable proportion of internet users are upfront about themselves. Nevertheless, it's still worth pausing to think about how your targeting controls might be unreliable.

Dressing for the occasion

Once we've honed our adverts in on the people most likely to become future customers, our targeting tools help us in a different way: They divide our targeted group into hyper-targeted subgroups which each receive marketing messages tailored just for them. The car manufacturer might determine that people in their 30s value safety and space for children above all else, whereas people in their 20s and are more interested in ecological footprint and electronic gadgetry. By emphasising the right benefits to the right groups, the manufacturer tells everyone exactly what they need to hear before they buy.

This sort of advertising segmentation is akin to owning multiple sets of clothes—for example, a conservative suit for job interviews and outrageous hipster attire for music festivals. Wear the wrong clothes to the wrong occasion and you won't have a good time.

The AND gotcha—avoiding unnecessary targeting

Because it's rare for a single targeting option to sufficiently hone in on the intended audience, we often layer options together so

as to combine their effects. On most platforms, each layered targeting option is related to the others with an implicit AND logical operator. If you aren't aware of what this implies, you risk putting a low ceiling over your potential reach.

Suppose a business wished to reach UK-based law students to sell textbooks written in English. They conjure up an image of their typical customer, and, based on this image, they set the following targeting options:

University Country	UK
University Degree	Law
Age	18-28
Language	English

Because these targeting options relate to one another with an AND operator, advertising will *only* be shown to people who simultaneously match all four criteria. But this high hurdle will overly restrict the reach of these adverts. Plenty of people in this business's *intended* audience (UK-based law students) may fall outside of the *targeted* audience, as defined by these targeting filters.

- Just because someone has their Facebook account set to a language other than English doesn't mean that they don't *speak* English. Indeed, the fact that a person studies law in the UK (as required by the other layers of the targeting) indicates that they almost definitely speak the language proficiently. The inclusion of this targeting option only served to exclude foreign students.
- Just because someone is over 28 doesn't mean that they don't study law. As it stands, this business is completely ignoring mature students—and these are the ones with the most money to spend.

This business would have been better off thinking about what criteria are implied by others, then throwing off whatever turned out to be superfluous.

www.dbooks.org

Let's look at another example. Our law textbook seller has been trying to advertise to law students on Facebook for months but is increasingly frustrated because many people leave their degree subjects blank. The seller then has the idea that anyone who likes the Facebook page of a university law society is probably also a law student, regardless of whether this person registered their degree with Facebook. The seller expresses this realisation through the following options:

University Degree	Law
Likes	An exhaustive list of law societies

For the same reasons as those given in the previous example, these targeting options leave the business short-changed in its reach. The implicit AND between the options limits advert reach to people who *both* registered law as their degree and *also* liked a law society. Because the intersection of these two groups contains fewer members than the group of people who just list law as their degree, this business's advertising reach would be even worse than it was before they decided to target law societies.

What the business really wanted was an OR between their targeting criteria such that anyone who EITHER studied law or liked a law society would be delivered the advert. Some platforms let you sandwich your targeting options with OR relations. When this possibility isn't present, you can achieve the same effect by splitting your adverts into two side-by-side ad groups that differ in their targeting options (one will target law society "likers" and the other, law students).

Oblique targeting

The most obvious targeting configurations are often oversubscribed. When pursuing these advertising avenues, expect stiff competition from other businesses with bigger budgets and deeper pockets than yours. In hotter markets, heavy competition may completely price new contenders out of the mainstream targeting options.

Dissatisfaction with typical targeting options can also occur in the total absence of competition. It's possible that you fully exhaust your leads from straightforward options, then find yourself in dire need of a fresh stream of customers.

This is the kind of environment that rewards the creative thinker. Oblique targeting is all about unearthing new pockets of leads in cobwebby corners or down grassy cul-de-sacs. If priced out of targeting law students directly (say by targeting university degree), how *else* can the marketer reach that same audience? This requires that the marketer imagine what it's like to be a member of that audience: What magazines, books, and blogs do law students read? What movies do they watch? What institutions are they members of? What tangential interests do they tend to have? Eventually the marketer will shout "Eureka" at having found some clever way to target by configuring some oddball combination of controls in their advertising cockpit. For example, the marketer realises that law students often engage in a special form of debating just for lawyers-in-training ("mooting"), and this realisation can be leveraged by targeting anything and everything related to it.

This kind of oblique targeting may not be as cost-effective as targeting directly, but it might nevertheless be profitable enough to be worth doing.

Negative targeting

Some sets are easier to define in terms of what they *don't* contain instead of what they *do* contain. When someone loses the "Q" tile in Scrabble, they don't complain that they only have "9 A tiles, 2 B tiles, 2 C tiles…" remaining. They simply say, "The Q is gone".

Negative targeting is a bit like this. Instead of defining who your intended audience are, you define who they are not. This is awfully useful.

For one, negative targeting is used to describe sets of customers that are impossible to describe with the available positive targeting options. If there is no option to target people who are "single",

www.dbooks.org

you could reach much of the same group by negatively targeting people who are "married". If there is no option to target people who don't watch TV, you could (theoretically) negatively target those that show no signs of interest in TV or TV shows.

In other situations, negative targeting can be used to ease the administrative entry of targeting configurations that would be awkward to describe in the positive. Without negative targeting, anyone who wants to target people living *outside* the USA would have to spend an hour in the advertising editor exhaustively listing every country other than the USA.

Another use-case for negative targeting is to fine-tune overeager positive targeting so as to better focus advertising budget on the underlying intended audience. In Google AdWords, negative keywords act like fussy editors who scrutinise and double-check targetings suggested by positive keywords. For example, a sweet shop owner who positively targets the (broad match) keyword "candy shop" will waste money on people looking up the 50 Cent rap song of the same name. By adding negative keywords related to the artist names or to the music domain (such as "50 Cent", "Olivia", "song", "listen", "youtube", "music", etc.), the sweet shop owner streamlines her targeting and saves money.

How are you to figure out what negative targeting options will fine-tune your audience? After all, not everyone is up to speed on their pop culture references…. A good general strategy here would be to start out with a purely positive targeting advertising campaign, leave it running for a while, then study the performance reports provided by your platform. Assuming the data is detailed enough, you'll notice that some subgroups have atrociously bad click-through rates; these you trim out of future targeting through negative targeting.[1]

[1]Given that Google AdWords is the most popular of all advertising platforms, I'll describe a specific workflow for gathering negative keywords here. AdWords compiles a report called the Search Terms Report that tells you what *actual search queries* triggered your adverts. This report will inform you that your bid on the keyword "candy shop" ended up displaying adverts 20 times for the query "50 Cent candy shop", 10 times for "download candy shop", and so on. This report also tells you the click-through rates and cost per click for each of these actual search queries. After leafing through it, you'll know what negative keywords you ought

Negative targeting isn't available on every platform yet; let's hope it gets rolled out more generally soon.

The edge

Keeping with the general societal trend of technological acceleration, advertising platforms are constantly outfitting themselves with brand new targeting possibilities. Not only this, but brand new platforms are forever arriving on the scene offering fresh opportunities for the bold. All of this novelty is a playground of potential for the intrepid advertiser quick on her feet.

Whenever a new targeting possibility is released to the world, it begins its life alone, massively underrated and undersubscribed by other marketers. This often means deliriously good deals for the nimble early adopters who show up at the great big auction house of online advertising. If you are quick to familiarise yourself with the latest batch of possibilities, you'll gain an edge—at least until your competitors crowd in after you.

In fact, it's important to stay up-to-date on targeting developments no matter what. Even if you arrive too late to the auction to avail of the bargain prices, it may still be worth your while to pay full price for use of a newly released targeting option that identifies your intended audience more finely than any previous technology could.

In addition to all this, there are special benefits accruing to advertisers willing to test the waters with newly launched platforms. The adverts shown on these platforms are just as new to the platform's regular users as they are to the advertisers paying for them. As a result, the platform's users haven't yet developed platform-specific ad-avoidance behaviour, which makes them (temporarily) more responsive to adverts. What's more, deep-pocketed corporate competition will be dissuaded from advertising during a platform's earlier days because it probably won't launch with the APIs and other tools that heavy buyers need and expect.

to set up.

www.dbooks.org

The inflationary effect of competition on advertising costs has been remarkable. During my first three years advertising with Google AdWords, I saw my cost per click rise by close to 300%. After this happened, I diverted much of my budget to Facebook advertising, which was then available at prices significantly less than what I was paying Google AdWords. Eventually though, Facebook advertising costs ballooned and stopped being the bargain they once were. Today, these advertising platforms are, at best, modestly profitable for my business, whereas they were once magnificently so. The moral of the story: The effort required to stumble through early, unpolished, and under-documented new advertising platforms is justified by the heavily reduced costs.

Synonym bombing

As mentioned earlier, much of the data that underlies targeting options originates from self-reported factoids given by users on the consumer-facing side of advertising platforms. When this data comes from options users select from pruned drop-down menus, this poses no problems. But complications arise if it comes from user fields that accept unrestricted input, such as text boxes.

Suppose that you are a seller of cat food and you wish to advertise on a platform where users list their likes in whatever ways they please. The relevant data, when re-appearing as a targeting option, won't be packaged into neatly standardised little semantic parcels. Some people will have liked "cat food"; others the differently worded "cat feeding"; or the more specific "food for Siamese kittens"; or a brand of cat food such as "WHISKAS"; or a blog intended for cat owners, like "The Reluctant Cat Owner's Journal"; or the idea of bargains, as in "discount cat food".

The infinite granularity of expression afforded by unrestricted input has its pros and cons for targeting. On the plus side, advertisers won't be pigeonholed into the advertising platform's preconceptions about what should be available to target. This gives advertisers the power to go after impossibly small niches (e.g., fans of Hodor from *Game of Thrones*), or to respond quickly to

novelty (e.g., a band that just made its first public appearance). On the down side though, granularity creates work. Advertisers won't maximise their reach unless they target the various ways their audience refers to whatever core concept is being targeted. I call this "synonym bombing". Although the technique isn't named this way elsewhere, it's still an officially sanctioned practice. For example, LinkedIn advises advertisers targeting by job title that, "If you're looking for an Account Manager, also try targeting a Senior Account Manager, Junior Account Manager, Director of Accounts, Account Strategist, and so on."

Undersubscribed targeting groups

Just because a targeting option perfectly fitted to your business *exists* doesn't mean that it'll be worth a sod for your business. Presence of a targeting *option* doesn't guarantee presence of any *targets*. The group may well be undersubscribed. For example, Twitter's behaviour targeting gives advertisers the power to display adverts to American homeowners whose properties are valued in certain ranges (e.g., "under $99,000", "above $500,000", etc.). The total number of Twitter users in every property value band *combined* comes to 13 million people. But Twitter has 65 million users in America.[2] Given that about 67% of Americans own homes[3], the math implies that Twitter is missing home value information for 4 in 5 of its users. For many advertisers, under-subscription of this degree will fatally dilute the targeting option's power. Cruder alternatives, like demographic and education information, may turn out to be more discerning.

The sum is sometimes more than the parts

As we saw in the last section, advertising platforms don't have complete information about their users. This problem occurs again and again—not all companies publish their workforce sizes on LinkedIn, nor do all Facebookers volunteer their relationship statuses. This leaves large gaps of unspecified data across many different targeting dimensions. Hopeful advertisers who target a

[2]http://www.statista.com/statistics/274564/monthly-active-twitter-users-in-the-united-states/

[3]https://www.wikiwand.com/en/Home-ownership_in_the_United_States

www.dbooks.org

specific company size on LinkedIn (10–100 employees) or relationship status on Facebook ("single") will therefore fail to reach people who are in reality a good match but cannot be identified through targeting. The fallout can be substantial—sometimes even the majority of platform members.

Many platforms have a special targeting setting for reaching the "unspecified"—those who fall between the gaps. The "unspecified" are worth advertising to if you have reason to believe that a satisfactory proportion of them are indeed part of your intended audience. For example, if you wish to reach unmarried people under 24, it may be worth targeting not just those whose relationship status is "single", but also those whose status is "unspecified". This is because it's reasonable to assume that young people are mostly single anyway, and that this fact will hold within the unspecified group—perhaps even especially so since those who are married are likely to broadcast this fact.

On the flip side, you are unlikely to have much luck efficiently reaching niche sub-groups (e.g., "open relationship") through targeting "unspecified". Those in open relationships only make up a tiny proportion of the underlying population, so you shouldn't expect them to appear with any worthwhile frequency in the "unspecified" group, at least not without the assistance of additional targeting layers.

Research advertising ease before committing

Before committing to any new business venture, there's great value in knowing whether you'll be able to meaningfully market through paid advertising. When available in sufficient volume and at tolerable prices, paid advertising makes launching a new business a thousand times easier.

Many advertising platforms give you estimates of how many people a given combination of targeting options ought to reach, and at what expected cost. By poking around in here, you get a feel for whether a proposed business will be easy, difficult, or impossible to advertise.

A Tour of Targeting Options

Let's take a guided tour of the main targeting options on the market right now. Once you're done with this section, you'll know your way around the modern territory as well as an old pro.

Geographical location

Most advertising platforms let you target countries, counties, cities, and postcodes; some even let you draw an N-kilometre radius around a chosen epicentre.

These geographic targeting features are most obviously of promise for businesses with a physical presence. There is a limit to how far people will travel to buy something, and advertising outside these limits is wasteful.

Not many people in small-town Spain will buy a guide to negotiating with a New York landlord, nor will a New Yorker buy a kitchen oven designed for the Spanish power system. Geographic targeting helps advertisers of region-bound products to restrict their adverts to the appropriate geographical bounds.

Residents of rich countries have larger disposable incomes that those of poorer countries. When you advertise to the whole world, you might discover that 90% of your advertising budget gets spent in India, yet you don't get a single sale there. This isn't a fictional example; it actually happened to me.

The concept of geographic location seems self-explanatory, but there's a hidden gotcha: It doesn't always make sense to classify someone as present in *one* location. Imagine a student who spends their weekdays at Oxford, their weekends with their partner in London, and their summer holidays in their hometown in China. No matter which of these three the student lists as their location, they will fail to capture the full reality.

Some platforms calculate the location from the user's IP address or network activity. This gives us a concept of "current location". But current location isn't necessarily the key considera-

www.dbooks.org

tion. While some businesses (e.g., local restaurants) *do* want to target people *currently* nearby, others are more interested in people who *live* in some place or other (e.g., house painters). The fact that someone is googling for painters while holidaying in France doesn't change the fact that they want their house in the UK painted.[4]

As we can see, which location is relevant depends upon the advertiser's specific motives. For this reason, it's sometimes best to completely remove location targeting whenever you've applied some other targeting dimension that implicitly restricts location. For example, if you wish to reach "law graduates in Oxford", it's enough to just target people who graduated any of the law universities in Oxford. An additional layer of location targeting would be superfluous.

As with all other targeting options, there's the potential for geographical targeting to help you reach your intended audience through oblique and convoluted means. For instance, Google AdWords doesn't let advertisers directly target employees of a particular company, but there's nothing stopping an advertiser from limiting their advertising to the smallest allowable radius around their target company's offices.

Time

Every online advertising platform targets by time in some form. The lowest common denominator of this is the power to manually switch a campaign on or off.

We most typically think of seasonality as something measured on the scale of winters and summers, and it's true that most of your efforts in targeting time will be spent adapting your adverts to these macro trends. If you sell fancy dress costumes, you'd be mad not to tailor your campaigns for Halloween, Christmas, and any of the other major holidays. If you sell ice cream, you'd be mad not to freeze your campaigns in winter.

[4] As if there weren't already enough ways to think about location, Facebook introduced the concept of "recently moved", which they define as people who recently changed the city listed in their profile.

Modern online advertising tools have brought us to the age of micro-seasonality, an age where the peaks and troughs of demand can be measured by day of week or even time of day. It would hardly be surprising if there were a correlation between night owls and potential customers of sleep-aid products. Neither would it be a creative stretch to suggest that work-related products would be more effectively advertised during office hours or that products intended for students would remain largely unseen on a Friday night.

How do you figure out these correlations? Through a combination of Analytics and detailed reports from previous advertising campaigns.

Gender

Some products are bought disproportionately by members of certain genders. To name a few: boxing shorts, tampons, bras, and protein powders. Targeting by gender can help you better zone in on your likely audiences. If you know that 97% of people buying your product are men, then advertising it to women doesn't make much business sense.

There's no need to go crazy with gender targeting though; in the presence of more discerning targeting options, it may be completely unnecessary. Even if an advertiser knows that it's mainly women who buy female birth control pills, it's patently worth advertising the pill whenever someone who isn't a woman expressly shows interest by googling prices for the pill. Maybe they are buying on behalf of their partner.

Only a handful of products are strictly for one gender or another; most merely show slight statistical differences in buying patterns which only justify proportional modifications to bid sizes for each gender (e.g., the advertiser might bump up bids 20% for women).

Sometimes you'll need to vary your advert creatives so as not to alienate one gender or the other. The seller of disposable razors should know that men are more likely to shave their faces and

www.dbooks.org

women more likely to shave their legs. Not tailoring their images and copy may hurt the retailer's conversion rates.

Age

For many of us, life unfolds in a predictable progression of stages. As kids, we are schooled. As young adults, we study or begin working. During this time, we move out from our family home and develop new interests. But with these new freedoms and abilities come new problems—we must now rent and maintain property, feed and insure ourselves, and earn enough money to support our lifestyles.

Age targeting is a crude targeting option that relies on these statistically typical paths through life. In a sense, its effectiveness is a barometer of cultural norms.

Age targeting is best used when combined with other options. Some time ago, I wanted to target undergraduate students. Back then, Facebook let you target members of a particular university but didn't give you any way to target level of study. (It added this feature a few years later). Despite this, I was able to reach undergraduates (skipping postgrads and staff) by restricting the advertisements to only students who were under 23. (Yes, I would have missed mature students, but this was the trade-off I chose to make.)

The typical users congregating on some platforms are a good generation older than those on others. For this reason, advertisers need to factor in the *implicit* age targeting that occurs on each platform. When I was a teenager, I used a social network named Bebo that was hugely popular in my home country, Ireland. I left for Facebook when I went to university, but Bebo continued for years afterwards to be thronged with early teens. By contrast, my godmother—nearly 30 years my senior—has only ever signed up to one social network: LinkedIn. As a woman working for a British corporation, her choice of social network is a professional necessity. Aged-related platform usage patterns go beyond social networks. My dad searches the internet with Bing because it was

the default search engine installed on his Windows machine, but no one reading this book would dream of using anything other than Google.

Device/computer/browsing technology

The advertisers most grateful for this targeting option will no doubt be the software vendors whose products are exclusively designed for some particular piece of technology. Why should they waste money advertising Mac-only applications to Windows users or Google Chrome plugins to people who browse with Safari?

Even if a software vendor sells a cross-platform product, these targeting options remain important for directing users to the landing page that matches each user's tech stack.

Depending on what device someone is browsing with, they are either more or less in the mood to purchase. Consequently, emphasising or de-emphasising adverts on certain devices (phones vs iPads vs laptops) can affect revenue in surprisingly large ways. At time of writing, customers are unlikely to complete long buying processes involving credit cards from their mobile phones, so many ecommerce businesses decrease their mobile advert bids. But other companies, such as call-out taxi services, would leave their mobile bids as high as their laptop bids, if not higher.

Interests (/ Likes)

Interest targeting is available in many different guises across the various platforms, so it's best to begin by sketching out the important strands which distinguish one manifestation from the next.

Recency

People's interests change over time, so each platform's way of categorising someone as "interested" matters a lot. Just because you were into mortgages six months ago doesn't mean you are today; once you've bought the house, you won't give financing another thought. Similarly, just because teenage-you liked a band on Facebook doesn't mean that your tastes as an adult haven't since matured.

www.dbooks.org

But your history of Facebook likes only gets longer. And the interests Google guesses you have (may) stay frozen in time.[5] Theoretically, the perfect advertising consumer would systematically prune their likes and meticulously register and deregister their interest elsewhere. But no one, at least no one remotely normal, does that.

As a result of this danger, advertisers targeting by interest should watch out for targeting algorithms that display adverts to people whose interests have grown stale. Your suspicions should be raised when the algorithm relies on data that could be months or even years old. By contrast, some algorithms are inherently current by design. For example, advertisers on Reddit who target by whatever subreddit someone is currently browsing are likely assured of engaged readers.

Private declarations

Some platforms harvest interest information from public user activity (e.g., timelines in Facebook and Twitter), whereas others pick up the info from private sources (e.g., browsing histories on Reddit or Google). This distinction matters because people skew their interests when they are in public.

For one, there's often an element of embarrassment or shame in expressing your true concerns. Not many people go to Facebook and like "cures for back pain", "ending heroin addiction", "how to meet the love of my life", or "pros and cons of telling religious parents you are gay". This means that the advertiser targeting by (public) interest will have a tough time reaching their intended audience. Instead, they will probably have to look for indirect routes, ones that rely on socially acceptable public activity trails. Perhaps the person with a bad back likes a brand of chairs known for their excellent spinal support. Or perhaps the recovering drug addict follows the author of a personal empowerment book that comes highly recommended within the community of recovering addicts but is rarely mentioned elsewhere.

[5] If you're curious about what interests Google thinks you have, check out their Ad Settings page: https://www.google.com/settings/u/0/ads/authenticated.

Social pressures and duties are another source of inaccuracy in publicly stated interests. People like things on Facebook (or tweet them) because of social pressure to conform, or social signalling, or the want to publicly support a friend's new venture which otherwise holds no special interest, or amusement at something said by a company (as opposed to appreciation of the goods they sell).

Granularity

In general, the more you can narrow down your targeted audience to match your intended audience, the better. The cat food retailer should prefer the advertising platform where they can target people interested in "cats" over the one that only lets them target people interested in "pets". This saves the retailer the wasteful expense of advertising to dog owners and fish owners, who won't have the slightest interest in cat food.

``Similar to''

Some platforms automatically calculate lists of people that they believe have similar interests to the people you are explicitly targeting. Frankly, as of early 2016, I have not once advertised with any platform's "similar audiences" feature without regretting it. But that's just my experience; you might have better luck.

Many platforms give advertisers more than one mechanism for targeting interests. By comparing these using the factors we just discussed, you'll be better equipped to choose the one best suited for your needs. For example, both Facebook and Twitter have concepts of *precise interests* and broader *predefined interest categories*. These differ only in granularity. Precise interests are ultra-granular (e.g., specific football teams, bands, etc.), whereas the predefined categories are more general (e.g., the concept of "football" or "rock music").

Let's close by looking at a sneaky little tactic: Advertisers who have well-known, established competitors can piggyback off their rivals' previous followings. The luxury handbag startup will target Prada lovers, just as the soft drink upstart will target Coca-Cola lovers. Be warned that this tactic may be illegal (check with a lawyer). And even if it is legal, it can sometimes backfire for an-

www.dbooks.org

other reason: Your competitors' fans may well be duds. I know of household brand names that hire shady companies to inflate their like and follow counts by using bots and outsourced like-farms.

Keywords

Famously, advertisers on Google Adwords can target based on search queries. The great thing about this targeting method is its to-the-second recency. Whenever someone googles "buy sandwich toasters", this, as you can imagine, is a damn good time to advertise that you're selling.

Less well known is Google's Display Network. Over two million websites have partnered up with Google (via their AdSense program) and agreed to rent out space on their pages in exchange for a cut of whatever Google earns by selling it to advertisers. Google is essentially a broker who automatically inserts ads whenever the on-page text of the partner's site matches a targeted keyword. The reach is phenomenal, with Google claiming that it encompasses 90% of all internet users.[6]

Keyword targeting isn't exclusively the domain of search giants. Twitter released their own keyword targeting based on the phrases recently tweeted or searched for through their platform. Should you happen to own a pizza delivery business, you could advertise to people who recently tweeted to their followers asking for a "pizza delivery" recommendation.

Keyword targeting should really be called keyword targetings. There is a whole family of differing "match types" which fine-tune when keywords count as being triggered. These keyword match types are (roughly) similar between the platforms, so I will focus on describing Google's.

Broad match

The most eager of all match types. If you don't set a match type when setting up your keyword targeting, Google sets it to broad

[6]https://support.google.com/adwords/answer/2404191?hl=en

match by default.

Broad match keywords match not just search queries that exactly match the original text entered, but also misspellings, different pluralisations, acronyms, abbreviations, synonyms, related searches, or searches that contain additional words. For example, targeting "sandwich grill" would cause the advert to display for

- sandwich grills - plural
- sandwich grilling - changed stem
- sandwich gril - misspelled
- grill sandwich - order changed
- buy sandwich grill - additional word at the start
- sandwich high-speed grill - additional words in the middle
- sandwich grill for sale - additional words at the end
- sandwich maker - synonym for one of the contained words
- panini grill - synonyms for both of the contained words

Modified broad match

This is a less extreme version of the broad match, the only difference being that synonyms and related searches are no longer targeted. But other than this change, modified broad match behaves just like broad match does. Out of the list of matching search queries we gave in the section for board match, everything except the following two would continue triggering adverts:

- sandwich maker
- panini grill

The motivation behind modified broad match is to give advertisers more control whenever synonym and related-search matching trigger too many irrelevant variants and waste profitability—which is basically all the time.

My recommendation to advertisers is to begin with broad match keywords for a few weeks and gather data about how well each synonym/related search performed. After this initial research period, create modified broad match keywords for all your top performers, then remove the original broad match keywords so as

www.dbooks.org

to prevent their wastefulness.

Phrase match

This match type is tighter than the modified broad match in that it requires that the search query contain the keywords as a single unit (i.e., in the same order and without words interspersed in the middle). But phrase match still matches search queries that have minor variations of the keywords (e.g., pluralisations, misspellings).

As such, the only difference between this and the modified broad match would be that now advertisements won't show for:

- grill sandwich (because the order changed)
- sandwich high-speed grill (Because a word was inserted in the *middle* of the phrase. By contrast, words added to the *start* or *end* wouldn't cause any problems because the unity of the phrase is still intact.)

Exact match

This is the tightest match type, but it isn't as tight as its name suggests. As with all the other match types, exact match still triggers advertisements whenever the search query contains the keyword with minor modifications (e.g., pluralisation, misspelling, altered stem). The difference between exact match and phrase match is that exact match has a zero-tolerance policy for additional words in the search query, no matter whether they appear before, after, or in the middle. This means that exact match won't display adverts for the following searches (which were fine with phrase match):

- *buy* sandwich grill (which showed ads with phrase match)
- sandwich grill *for sale* (which showed ads with phrase match)
- grill sandwich (which did not show ads with phrase match)
- sandwich high-speed grill (which did not show ads with phrase match either)

Remember that there is also negative keyword matching that the marketer can use to prevent adverts from displaying. For exam-

ple, if the retailer of a sandwich maker knew that a model could not prepare paninis, then she might negatively target the keyword "panini" while positively targeting broad match "sandwich grill". This means her advertisements wouldn't show for "panini grill", but would for other synonyms like "toastie grill".

Just as with interest targeting above, keyword targeting also opens up the possibility for you to piggyback off your competitors' brand names/trademarks in the hopes of convincing their leads and customers to switch. Google specifically states that they will not investigate or restrict the use of trademark terms as keyword targets, even if they receive a trademark complaint.[7] Despite Google's hands-off approach, this practice is still in a legal grey area; you risk legal consequences if you choose to pursue it.

Your fans

Plenty of businesses have a social media presence which they use to build a following. Their efforts are motivated by the promise of free marketing—the possibility to reach all their hard-won fans without paying a further dime. In the early days of social media, this actually worked; marketers experienced a golden age where almost anything they posted on social media would reach their fans.

This was not to last. Nowadays, marketers would be lucky if even 10% of their fans saw their social media posts; such is the jump in competition for a user's attention. It's commonplace for Facebook users to have friend counts in the thousands, each of whom posts more content than ever thanks to the ubiquity of share buttons, powerful new mobile devices, and the continuing myth that photographs of cappuccinos are of any interest to anyone. What's more, every Facebook member has been accumulating an ever-growing mass of group and page memberships for as long as they've been part of the network. Each of these pages continue churning out content in the struggle for attention, so now, instead of competing against 10 other pages, you're competing against 100.

[7] https://support.google.com/adwordspolicy/answer/6118?hl=en

www.dbooks.org

All profitable businesses are born out of scarcity, and this lesson hasn't been lost on companies like Facebook, who are now capitalising on this increasingly hot commodity. In the manner of a trendy New York nightclub, advertisers can now tip the doorman and skip ahead of the penniless multitude vying for a chance to be seen. One such bribe is Facebook's option to "boost" a page's post, which basically gives a brand's social media the same airtime as it would have gotten before the competition got steep. Or alternatively, you could target by "connections", whereby businesses can target their regular adverts to existing fans of their company's own pages, apps, and events.[8]

Connection targeting isn't just limited to reaching people who liked your page/app/event. It optionally extends to their *friends*. This feature is worth enabling when you believe it likely that the interest you're targeting will be shared by a *significant* proportion of a person's (probably very extensive) social network circle. (Remember, this circle probably includes their family, childhood friends, university friends, one-night stands, work colleagues, etc.). Just because someone who is otherwise civilised has a shameful obsession with the *Jersey Shore* TV show doesn't mean that more than a pair of their Facebook friends have the remotest interest in such things. On the other hand, I can imagine situations where the overlap of interests would be strong. For example, the average computer science student is likely to have a generous helping of friends they've met on their course who share similar interests.

Idealistic advertisers may scoff at the thought of paying to reach fans they had already won over organically. It feels desperate and dirty, like paying your brother to come to your own birthday party. But consider the pragmatic realities: Your fans—assuming they were legitimately earned—know and love your brand. They are the people in the world most likely to buy from you. Sure, you already invested heavily in winning their initial affection. But now that they love you, the marginal cost of converting them into paying customers is going to be far lower than the cost for

[8]It's not possible to employ "Connections" to target fans of pages *others* own.

fresh leads. Therefore, spending a little extra money on reaching these existing fans is not only rational, but optimal.

There's another huge benefit to connection targeting: It leverages social proof. When you notice that your buddies are going gaga over an advert you'd otherwise ignore, that's going to grab your attention and encourage you to give it a look. As such, Facebook is wise to slap a Like button on their adverts and highlight which friends appreciated it.

Other people's fans

Twitter has a wonderful targeting option that lets advertisers target the followers of *anyone* on Twitter—be they competitors, enemies, celebrities, bots, politicians, or your company's CEO. Other people's Twitter accounts now become user lists for you to advertise to, opening up endless possibilities for creative targeting. This is the marketing equivalent of picking up a Mario Kart power-up.

The most obvious use-case would be targeting the followers of a competitor, thereby poaching some of their hard-earned followers. But even if you have no direct competitor to piggyback off of, you could achieve much the same advantage by targeting businesses that serve different niches of the same market (e.g., a fitness supplement manufacturer might target followers of a chain of gyms).

Another use-case, this one more for the branding-conscious advertiser, would be to target the followers of a thought leader relevant to their customers. For example, a trendy shoe designer might advertise their pairs to the followers of a street fashion guru.

Uploaded lists of customers/leads

Advertisement isn't just about finding completely new customers—it's also about reaching out to existing ones and enticing them with new offers.

Facebook, Google, and Twitter let advertisers upload lists of

www.dbooks.org

email addresses then target these people through their platforms. The obvious limitation is that the advertising platform has to have someone in its system who matches that email address. The key metric here is the "match rate", and this varies quite a bit by platform. In 2015, Larry Kim did an experiment with 350,000 email addresses and found that Google and Facebook both matched about 50% of them, whereas Twitter only matched about 10%.[9] Regardless of which platform you're using, you'll be able to nudge the match rate upwards if you keep all your customers' old email addresses around instead of deleting them.

Businesses without email addresses on file for their customers won't necessarily miss out on this party. Facebook lets advertisers upload other types of identifying information, such as phone numbers, Facebook user IDs, and mobile advertiser IDs.[10]

Even when emails are collected by a business, these alternative identification methods are still worth collecting because they help bump up the match rate even further.

Leaving aside potential regional-specific issues with data privacy laws and website policies, there's nothing technically stopping an advertiser from procuring a list of customers from a third-party source. Two businesses wishing to reach similar audiences might thus agree to share their uploadable customer lists to mutual advantage.

Many business transactions still occur offline (e.g., in high street stores). This creates a lacuna for marketers wishing to advertise to their previous customers online. Luckily, uploaded customer lists can assist in smoothing over this problem. The trick is to unite a customer's online and offline identities through something that both existences share. This something is usually an email address. Having gathered the email address at the store, the customer can then be targeted with follow-up online advertising.

[9] http://www.wordstream.com/blog/ws/2015/10/05/adwords-customer-match

[10] Mobile advertiser IDs are basically unique device identifiers for mobile phones. Apps can be programmed to detect these IDs and shoot them off to the marketing department.

Uploaded customer lists also play a role in negative targeting. For whatever reasons, some customers will be unlikely to buy again. This could happen, for example, if your website is dedicated to selling just a single book; once someone has bought it, what's the point in continuing to advertise to them?

Finally, some platforms have features to find "similar audiences". This is closely related to a targeting mechanism we already looked at, that of similar *interests*.

Snipering

So far, we've spoken as if uploading lists of previous customers is the only way to leverage uploaded list targeting technologies, but this isn't quite true. Advertisers who so desire can alternatively choose to compile lists of likely leads using publicly available information. For example, an advertiser could research a list of Twitter handles they wish to advertise to (e.g., @jackkinsella, @WhiteHouse), then upload this list through Twitter's Tailored Audiences.

This is the online advertising equivalent of cold-emailing, and will no doubt be of great value to salespersons wishing to get their messages in front of decision makers they've already identified.

LinkedIn also has a way to achieve a similar effect, albeit on a smaller and more labour-intensive scale. It would be a stretch to call this advertising *pure* because it doesn't involve regular advert creatives with pictures, headlines, and all that jazz. Instead, the marketer sends out a special (paid) LinkedIn email (called an "InMail") to any LinkedIn member they want, even if the marketer doesn't know that person or have any connections with them. This tactic is very useful to the recruiter struggling to fill a position or the accounting software salesperson hoping to reach a company's CFO. If the CFO has no publicly available contact details elsewhere on the internet, InMail may be the only way to get through to her directly. This tactic risks bordering on spammy, so stay tasteful in your message and keep in mind any potential legal consequences.

www.dbooks.org

Behaviour

Behaviour targeting is a messy hodgepodge of next-gen targeting options based not on information someone keys in about themselves on some social network or other, but rather their *actions*, as reported by offline sources like customer loyalty programs or online sources like purchasing activity or subscriptions. Here's an example (from Twitter's documentation) about how behaviour targeting has been applied:[11]

> *Nestle's [@Butterfinger](https://twitter.com/Butterfinger) leveraged partner audiences for the launch of its new product, Butterfinger Cup Minis. By targeting Promoted Tweets to Twitter users who frequently purchase peanut butter candy in grocery stores, convenience stores and other in-store destinations, [@Butterfinger](https://twitter.com/Butterfinger) achieved a 52% lift in engagement rate compared to Nestle's overall performance in 2014.*

Behaviour targeting is available on Twitter and Facebook at present, although they don't gather the data themselves—this comes through partner companies. At the moment, behaviour targeting is undersubscribed (i.e., data is only available for a fraction of users, so the reach won't be what it should be). But as commercial transactions are increasingly digitalised in coming years, this number is sure to rise.

There are far too many different behavioural targeting options to list here, so I'll just mention some of the most interesting examples on each platform. Be warned too that the available options depend on what country you are advertising into. This is because data availability through partner companies differs from country to country. Places with strong data-privacy rules, such as Germany, won't have half the options available as a place with laxer rules, such as America.

Twitter

- Creditworthiness

[11] https://blog.twitter.com/2015/introducing-partner-audiences

- Yearly income
- Home value
- Liquid assets
- Investment ownership
- Net worth
- Shoppers who spend in certain retail categories (e.g., "fitness" or "electronics")
- Shoppers who buy a particular brand (e.g., Reebok or Vans)
- Donation habits (e.g., "animal charities" or "political liberal causes")
- Cat or dog ownership
- Political party affiliation (US only...this one must be great for raising donations!)
- Automobile ownership (e.g., how old someone's car is and what brand they own)
- Sporting interests (e.g., people likely to attend live soccer games, go camping, or play tennis)

Facebook

- Expats living in a specific country (e.g., in Germany or China)
- Travelling behaviour (e.g., frequent travellers, business travellers, commuters, users of travel apps, people currently travelling)
- Facebook specific digital activities (e.g., people with histories of making Facebook payments, creating events, admiring Facebook pages, or playing Facebook games)
- Other (non-Facebook) online activities (e.g., players of console or browser games, uploaders of photos)

YouTube

Through Google's Display Network, businesses can direct their advertisements to appear on specific YouTube videos or channels. This exciting targeting option is useful for a range of marketing purposes such as promoting your YouTube channel to viewers of a rival channel, advertising your music instrument ecommerce store in videos that talk about synthesisers you stock, or promoting your jazz school on videos dedicated to teaching piano im-

www.dbooks.org

provisation through video.

YouTube also has a video remarketing feature that enables video uploaders to target people who've "interacted" with their videos in the past. Possible interactions include watching videos, taking an "action" on a video (like, dislike, comment, share, etc.), or subscribing to the YouTube channel.

TV

Whenever a TV program airs its latest episode, fans buzz about it all over the web. From the flurry of activity on its own platform, Twitter can easily tell which of its users are engaged with which TV programs. Twitter then harvests all this info and packages it up into neat little TV targeting options.

TV targeting comes in two variants. The first, *new airing* targeting, restricts advertisements to displaying around a show's airing window. New airing targeting must be contrasted with *continuous* targeting, which instead displays the promotional messages at any time during the advertising campaign—during TV show airing or otherwise.

TV targeting will be most useful to businesses already buying advertising space on TV and wishing to supplement their messages there with Twitter campaigns. I imagine too that businesses selling merchandise related to specific TV programs will also want to tinker with TV targeting.

Major world event

Whenever some newsworthy event is happening—a major conference, an election, or a movie release—you can be sure that people will talk about it on Twitter.

Unsurprisingly, advertisers want to insert themselves into the middle of these conversations. Even for companies that don't sell products directly related to the event, timely and relevant wit can be remarkably effective at winning over the affections of the general public. For instance, during Super Bowl XLVII, a blackout struck and Oreo quickly quipped, "You can still dunk in the

dark". This tweet, despite being promotional, was famously well received.

Twitter has a special targeting option just for these moments. Marketers are given a pre-populated calendar of major upcoming events around which they can schedule their advertisements. Behind the scenes, this is powered primarily by tweet consumption and engagement data.

This feature appears to still be in an early state of development. As I write this, the number of events on offer is only a few hundred. This rather paltry selection is organised into the following categories

- Conferences (e.g., Microsoft BUILD, ad:tech)
- Sporting events (e.g., the Boston Marathon, the FA Cup Final, the Miami Open)
- Entertainment events (e.g., the Billboard Music Awards, the Cannes Film Festival)
- Movie releases (e.g., the latest X-Men movie)
- Political events (e.g., Super Tuesday)
- National holidays (e.g., St. Patrick's Day, Mother's Day)

Relationships and family

Targeting by relationship status will most obviously be important to advertisers in the dating space. Facebook lets advertisers target not only the classic denominations (e.g., "single" or "married"), but also a selection of narrower options including "divorced", "open relationship", and "long distance relationship". These could come in handy for advertisers of niche services, such as software that facilitates couples watching synchronised Netflix videos together from afar.

If you want to create appropriately enticing creatives for dating services, you'll want to know your audience's sexual preferences. To this end, Facebook lets advertisers target by whichever sex a person registers they are "interested in".

Freshly engaged couples often publicise their plans to marry

www.dbooks.org

on Facebook, and, as is to be expected, Facebook capitalises on this juicy titbit by letting advertisers target by engagement status. Wedding planners, photographers, and catering companies love this.

One thing leads to another in relationships, and sometimes that other thing is actually a person. Through Facebook's knowledge of its members' family relationships, it is able to provide advertisers with the option of targeting to parents who have children in certain age ranges (e.g., "parents of preschoolers", "parents of pre-teens", or "parents of teenagers".

Users of your own app

App developers often make their money by convincing users of their free app to upgrade to a paid version or complete an in-app purchase. For this reason, it's sometimes worth their while to advertise to their *existing* app users, even when these users have logged out of the app and are hanging out somewhere else on the internet. This functionality is no pipe dream—it's available on both Google and Facebook.[12]

Let's talk specifically about Facebook's feature set, it having more developed functionality that Google. Before using it, you'll need to do a mandatory bit of integration work to ready your app and establish a connection between it and Facebook. Once done, you are already able to advertise to your app users. But this might not be good enough, because all the users will be clumped together into an undifferentiated whole.

You can further refine your targeting by pinging Facebook with in-app events. These come in two varieties: custom and predefined. The predefined ones are designed to cover the basics, such as:

- App install
- App launch

[12] Google offers this feature as part of its remarketing toolkit; Facebook as part of its custom audiences.

- Search
- Add to cart
- View item
- Purchase
- Complete registration
- Complete tutorial
- Level achieved
- Achievement unlocked
- Rated

Most event notifications also accept additional parameter(s). For example, the "search" event could also send along the query string; the "add to cart" event could lead to the cart value, line item, and currency; the "achievement unlocked" event could result in a description of the exact achievement. You can then filter by this information to target the most profitable leads.

If needs be, you can also create your own *custom* in-app events that capture whatever is most important to your business for segmenting and targeting users.

Through Facebook's lookalike audiences feature, you can further your advertising reach beyond your original lists of app users. Facebook's lookalike audiences are supposedly quite good when fed app user lists (one company claimed that it increased their app's revenue six-fold). But of course, your mileage may vary.

There are some apps that only work when groups of friends have all downloaded it (e.g., messaging apps, work collaboration apps, or multiplayer games). These app developers will love Facebook's option to target the friends of whomever has installed an app.

Users of other people's apps

If you are launching a new app, there is no better group to advertise to than users of competing or similar apps.

Although this might seem hard to believe, Twitter collects and maintains a list of all the apps that its iOS and Android users have installed on their phones. They dress up this cheeky data grab

www.dbooks.org

as the "app graph" and use its contents to power their "installed app" targeting option. This lets advertisers target users who install apps that fall into certain predefined categories (e.g., business apps, finance apps, health & fitness apps, or reference apps).

Google has a different approach. Instead of only displaying ads to users of its own websites/apps (as Twitter does), Google syndicates its adverts to users of *other apps*, so long as those other apps are in cahoots with Google through its expansive Display Network. Impressively, Google lets advertisers target named apps (or "placements", in their lingo). At last, you can rest assured that the fine users of Extreme Fart Recorder PRO will hear your promotional messages! Because specific app targeting happens through Google's Display Network, these ads can be configured to appear on apps running on *any* platform (including iOS). Unfortunately, there's a big limitation to Google's otherwise compelling app targeting: The ads only show whenever a user is currently within that app, which may not be often at all.

Not everyone is willing to spend money paying for apps. In recognition of this, Google has added a targeting option that limits advertisements to only display to users who have previously purchased a paid app or made an in-app purchase from the Google Play Store.[13]

In addition to targeting specific named apps, advertisers on Google can also target broad (predefined) app categories, these being the same as those used on its app store.

Visitors of your own website (remarketing)

There are good reasons to advertise to people who've already visited your website.

First is the possibility of segmenting your audience so as to target only the visitors who have the highest likelihood of buying. This is done by filtering your visitors based on some action they performed on your website, such as adding a product to their cart or

[13]https://plus.google.com/u/0/+GoogleAds/posts/Wj6AjvvsW8d

spending more than 20 minutes browsing. If your remarketing list consists of only the hottest of leads, you can confidently bid more per click when you need to beat out a competitor in the auction for advertising space.[14]

Second is the idea that advertisements only peak in effectiveness after being seen multiple times. By continuing to deliver your messages after someone has left your website, you bolster your brand and help transform mild interest into a compulsion to buy.

Third is the trick of making your business seem larger than it really is. By showing follow-up adverts on Facebook, Twitter, and (especially) the zillions of websites in Google's Display Network, you create the impression of omnipresence. This inspires trust in conservative old customers who are scared of dealing with brands they haven't heard of. (As internet users become more and more accustomed to remarketing, this effect will probably wear off.)

Configuration

The first step in getting started with remarketing is to integrate the advertising platform's JS snippet/remarketing pixel into your website. Next is to update your privacy policy to alert visitors that they will now be subject to remarketing. This is necessary because remarketing has strong data privacy implications. Once that's done, create skeleton remarketing lists as soon as possible, ideally at website launch. Delay in setting up these lists may leave you with significantly less leads to remarket to than you would have had if you'd gotten your act together earlier.[15] The last bit of remarketing configuration consists of choosing the trigger conditions that will cause a visitor to be added to your lists. The usual default trigger is someone visiting a particular URL (e.g., the URL

[14] But bidding more per click when there is no threatening competition is unnecessary and will increase your cost of advertising over the long term. Therefore, it's important to know whether you are losing a share of available impressions to your competitors. In Google AdWords, the extent to which this is or isn't happening is summed up in a statistic called "impression share". If you already have an impression share of 90–100%, there's no point in raising your current bids.

[15] "*May* leave you with significantly less leads" because *some* types of remarketing lists on *some* platforms can be populated retrospectively. That said, most cannot, so it's safer to create skeleton lists as soon as possible.

www.dbooks.org

"/checkout"). Other possible trigger mechanisms are the visitor having spent so-and-so long on your website, or the presence of special custom variables which your website sends the advertising platform (e.g., you might have a "frequent-visitor" variable that gets sent whenever someone has been to your website five times in the last month).

It's important to mention that you can choose how long someone stays on each remarketing list. Lists can have short or long stay times. Short times signal recent interest/engagement and are useful to advertisers who want to strike while the iron is hot. Longer stay times, by contrast, are more like capturing email addresses. These lists are useful for customer-lifetime marketing strategies, such as announcing sales or launching new services of interest to old customers. Often it's worth having parallel short- and long-stay lists, with each list serving a different marketing end (e.g., you'd have one list for people that visited the checkout page in the last seven days and another for everyone that ever visited it).

MUST-HAVE STARTER LISTS:

- **All visitors** - A handy catch-all list useful for general advertisements (e.g., "sale now on", "introducing our newest service").
- **Cart abandoners** - Contains people who were on the verge of buying something but didn't follow through. It's worth advertising to these folk because you have a good chance of convincing them to go all the way and complete that purchase—especially if your remarketing campaign offers them a little discount.
- **Previous customers** - In many businesses, past customers are likely to return and buy again. Ideally, the adverts you show them should emphasise products that are often bought together with whatever that past customer already ordered.
- **Lurkers** - People who spent a certain minimum amount of time on your website—say "2 minutes" or "6 minutes". This is an indicator that they are seriously considering your offer.
- **Churners** – People who used to subscribe to your service but have since cancelled. When you release new features or

reduce your pricing, you might be able to convince these people to start their subscriptions again.
- **Window shoppers** - People who viewed a particular product (or category page). This is really a bunch of lists, one for each major product or category you stock. Having these lists around will enable you to segment your leads and only remarket products that each lead previously showed interest in.

It would take ungodly amounts of time for companies with expansive product ranges to set up hundreds of thousands of remarketing lists so as to segment their leads properly. To this end, Google created something called Dynamic Remarketing. If you've ever visited Amazon and looked at a Bose speaker only to see the same model advertised a minute later on another website, then you've been exposed to this technology. Dynamic Remarketing works by pulling product data from a feed on the advertiser's website and automatically generating remarketing ads.

Remarketing overload

There is a risk that advertisers employing remarketing end up bombarding their audiences with adverts everywhere they go. This can come across as annoying, or even stalkerish. In gauging how pushy you're being with your remarketing efforts, remember to consider the *total effect* of your remarketing across all platforms. If a customer sees your remarketing ads through Facebook, Google, and also Twitter, their exposure to your advertisements is tripled....

One way to combat this effect is with frequency capping. This works to limit the number of ad impressions shown to one person in one day. Personally, I configure each platform to show only one impression per day.

The second way to soften the potentially pushy vibes from remarketing is to limit how many days your campaign runs for. This can be controlled most easily by advertising to suitably short-stay remarketing lists (e.g., seven days). After this time is up, any given individual on that list would no longer see the advert (unless they

www.dbooks.org

visited your website again…). The advertiser without the benefit of short-stay lists can exert cruder control by simply turning the whole remarketing campaign off after a number of days.

Negative remarketing

There are also *negative* remarketing lists. These can be combined with other targeting options (or even positive remarketing lists) to further narrow the targeting of a campaign. The classic use-case would be companies selling a single product. Once someone has bought this item, there is no chance of them buying again; therefore, it makes no sense to continue advertising to them.

Visitors of other websites

Targeting visitors of *other* websites is perhaps the most basic of all the targeting mechanisms ever built, one that is as old as the internet itself. Even to this day, it's possible to directly buy advertising inventory from any website offering it. That said, inventory is nowadays mostly sold through bigger brokers, with Google being the biggest of them all. Through their Display Network, you can rent space on any of the gazillion websites signed up to Google's AdSense program. Indeed, thanks to their "automated placements" feature, you needn't choose the exact websites yourself; instead, you can rely on other targeting options (e.g., keyword matching) to pick out appropriate websites.[16] After leaving this running awhile, you can swoop in and promote the most profitable placements, and shoo away the bums. If this sounds a little familiar, that's because it is: We already encountered "placements" in the section about targeting users of other people's apps.

It's possible to go even deeper than choosing individual websites in your manual optimisations. Websites today are rather diverse affairs, and perhaps only a fraction of their pages are relevant to your business. One way around this waste is to ask Google to limit the targeting to named folders (or specified URLs) of the placement website. If you are a seller of digital cameras, there's little point in indiscriminately advertising on *every* page of a review

[16] https://support.google.com/adwords/answer/2580292?hl=en

portal—but it wouldn't be a bad idea to advertise on the sections pertaining to photography.

Of course, pruning placement folders may quickly become tedious, so you may prefer the automated option of combining placement targeting with keyword targeting. Now, advertisements will only be shown on pages that also match for keyword content.

Because webmail clients count as websites, it is possible to shoehorn your adverts into email inboxes. To advertise in Gmail, set up a placement for *mail.google.com* then create some special Gmail Ads from their Ad Gallery.[17] Once approved, your ads take an expandable, email-like form and show up under Gmail's "Promotions" tab.

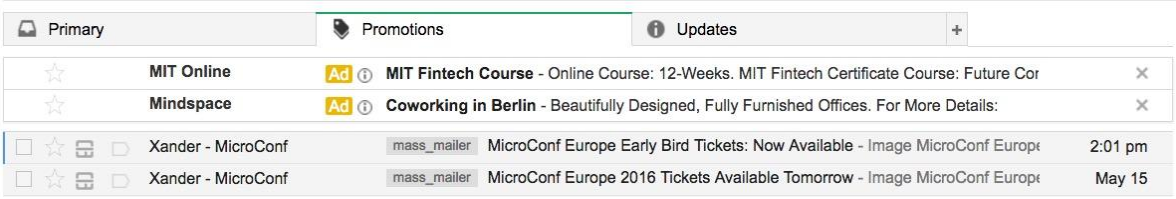


Figure 19.1 – *Source: Google.com*

The websites displaying the adverts in the Display Network come from all walks of life: extreme right-wing newspapers, obituary listings, parked domains, lobby groups for legalising euthanasia or banning abortion, etc. Needless to say, juxtaposing your brand with these websites may be undesirable. Exactly how undesirable depends on how sensitive or family-friendly your brand is intended to be. Edgier brands, such as food delivery company Eat24, actively seek out cheap advertising options on adult/-porn advertising networks, buying up advertising real estate that other brands wouldn't touch with a 10-foot pole. Another important consideration is your target demographic: Young internet users are usually aware that there's an element of randomness to where advertisements show up, and know that an advert's presence on an undesirable website doesn't imply that the advertiser

[17]There are many gotchas; see this video walkthrough from Google: https://support.google.com/adwords/answer/6105478?hl=en.

www.dbooks.org

was aware of or approving of the connection. The same cannot be said for older, less technologically savvy web users.

Google offers various options for taming the range of placement websites. The most surgical and precise of these is blocking specific unwanted sites through negative placements. Also rather precise (and tedious) is the option of blocking websites that contain certain "negative keywords". For those who don't want to fiddle endlessly with these options, there is an easier general solution: Google has high-level filters available in their "site category options" which let you disable anything from military content to advertisements placed on error pages. The current list is printed below:

I'd like to finish with some cautionary words about the technical format of placements: There is a meaningful distinction between domain names with and without the preceding "www". Google's documentation tells us that excluding example.com is general enough to also exclude www.example.com, but excluding www.example.com means your adverts will still appear on example.com.[18] Because many websites respond to both "www" and the bare domain name, your placements (or exclusionary placements) could end up far less comprehensive than intended.

Language

Given that most platforms offering language targeting also offer country targeting, this feature only has limited applicability. One such application is tailoring the ad copy in countries where multiple languages are spoken. For instance, an advertiser in Belgium would show ads either in Dutch or French, depending on the primary language of whoever's on the receiving end.

Language targeting also shines when targeting subpopulations within a country that are identifiable through their speaking a language other than the national norm. Someone residing in Spain but speaking only English is probably an expat, and is therefore a good target for adverts for Spanish courses or English-speaking

[18] https://support.google.com/adwords/answer/2454012?hl=en

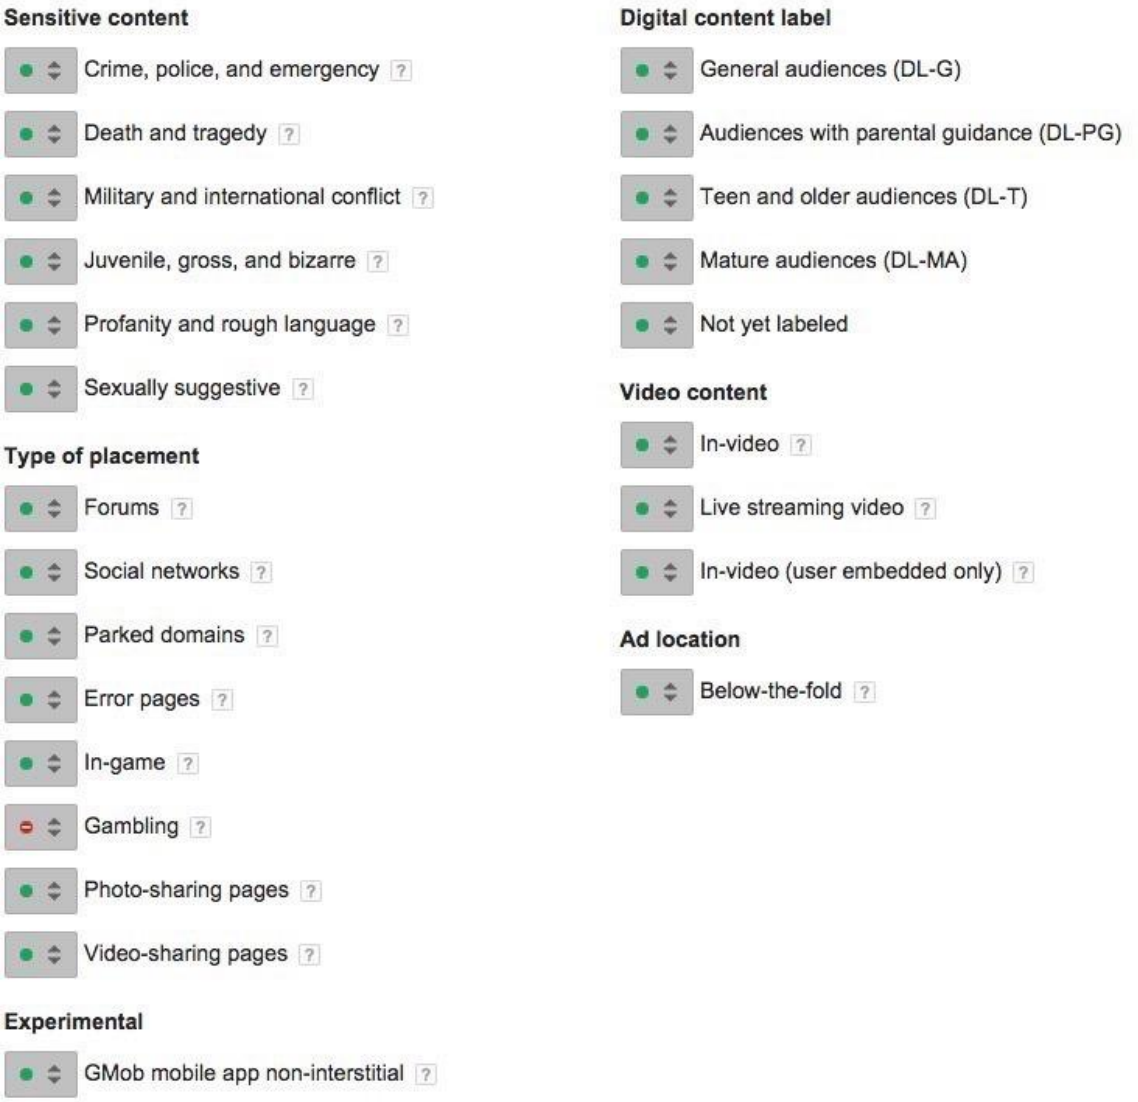


Figure 19.2 – *Source: Google.com/adwords*

www.dbooks.org

tax advisers.

Education

Facebook and LinkedIn let you target those who have attended a given educational institution, studied a specific field, or attained a certain degree. This sometimes extends to negative targeting. LinkedIn gives the example of a degree-awarding education provider *negatively* targeting people who already have postgraduate degrees.

Facebook has a few especially nifty education targeting additions. For one, you can specify a range of undergraduate years, so that adverts only show to one particular cohort. This enables you to run the same campaign every year, but to different cohorts, excluding from targeting those who saw your adverts in the previous academic year. Facebook also lets advertisers target by general education level, such as "at high school", "at university", "graduate", "doctorate degree", etc.

As you'd expect, education targeting is of most use to those in the business of education: MBA providers, online course purveyors, textbook publishers, or tutoring services. Another major consumer of these features would likely be recruiters who hope to snap up fresh graduates for their clients.

Workplace and job targeting

The primary purpose of these targeting options is to reach *companies*. Because companies are nothing but collections of individuals, the B2B advertiser has only to reach *employees* at a particular company in order to reach that company itself.

There are, of course, cheekier use-cases for workplace targeting. I know of a PR consultant who saturates employees of relevant trade journals and magazines with messages about his clients, thereby subconsciously kindling these publications' interest.

LinkedIn, as you would expect, does workplace targeting best, not least because its members have profile info that is more complete and up-to-date than anything on offer elsewhere.

LinkedIn (along with Facebook) lets you target employees of a particular company (e.g., Microsoft or Nike). This is a godsend once you know which companies you want to pitch to or partner with. Both advertising platforms also offer targeting by job title and industry. Targeting by job title is fully granular, so you have to synonym bomb a list of variations in the hope of matching the myriad self-reported job titles on file. Targeting by industry isn't granular at all; instead, you target broad predefined categories like "Healthcare", "Education", or "Tobacco".

LinkedIn (alone) lets you target by job seniority (e.g., Entry, Manager, Director, CXO, Partner) and job function (e.g., Accounting, Legal, Marketing, Operations). These options let you hone in on the decision-makers working at the relevant parts of the company you're pitching too. There wouldn't be much point in your financial services company advertising to your potential client's kitchen staff.

Sometimes you want to advertise to only the biggest players, and sometimes you only want to reach the little guy. LinkedIn lets you do both by means of its targeting by company size. Through this option you can choose to advertise to companies with only one employee all the way up to behemoths with over 10,000.

LinkedIn members can form "groups" for whatever purposes they like (e.g., local networking, professional associations, topic-based support). Unsurprisingly, LinkedIn lets you target these. To get some ideas about *which* groups to target, check out LinkedIn's groups directory and search.

Recruiters, professional training course providers, and so on will be fans of LinkedIn's targeting by skills. If you're organising a conference for a certain technology (e.g., Rails or CSS), then this is the targeting option for you.

www.dbooks.org

20 | Paid Advertising: Ad Creatives

We already explored the general principles of copywriting in a previous chapter. The chapter you are now reading is really just an extension and qualification of those points for the narrow case of *off-site* advertisements (e.g., adverts you place on Google AdWords or Facebook).

Image Choice

Not all advertising platforms let you put an image in your advert (e.g., Google's Search ads are text only). That said, whenever an image is allowed, it is, without a doubt, the most important part of your entire advertising creative.

The primary purpose of the image in an off-site advertising is to grab someone's attention. Through years of experimentation, I have found that this can be best achieved by using some combination of the following:

- Visual novelty/things that look weird (e.g., strange shapes or designs, attention-grabbing colours, anthropomorphic food, cupcakes as big as a house, 20-meter-long eels, photos with unnatural amounts of brightness or contrast).
- Danger (e.g., tigers, fire).
- Food (e.g., delicious cupcakes, fancy coffees).
- Sex (e.g., beautiful people, suggestive poses, vegetables or household tools arranged into lewd formations).
- "The guillotine": a photo of a person looking up out of the page, such that they seem to be looking right at the reader. (The typical Facebook profile photo would count as a "guillotine").[1] The apparent reason why this technique works is that it is difficult to ignore someone looking directly at you, even if that someone is only present in photographic form.

[1]This technique appears to have originated in *Ca$hvertising*–http://www.cashvertising.com/.

- Identification (i.e., the person depicted in the photo is of a similar age to the target audience).
- Amateurish photos. These work because people tend to gloss over anything that looks too much like a "stock" photo. People's attention will still be drawn by photos that are shoddier looking.
- Photos with heavy contrast, either within the photo itself or between the photo and the background where it appears (e.g., the photo's contrast with the greys, blues, and whites of Facebook).



Figure 20.1 – *Source: Pixabay.com*



Figure 20.2 – *Source: Pixabay.com*

www.dbooks.org



Figure 20.3 – *Source: Pixabay.com*



Figure 20.4 – *Source: Pixabay.com*

Are there any overarching theories that explain why these various photo styles work? One theory is that many of these categories, such as sex and danger, have a basic evolutionary significance to our reptile brains, and so they are more salient in our consciousness. Another theory is that some of those photo categories stand out in people's consciousness simply because they contain novelty, and our brains have a penchant for the new. Yet another theory explains these photo's effectiveness from a mercilessly mechanical perspective: Our eyes and brains are imperfect machines, and therefore images that are easier to discern (say because of colouration or lighting or some other property) will be more likely to register with us.

Contrary to most people's expectations and natural inclinations, no conceptual connection is necessary between your advertising image and the product you sell. Remember, the goal of the image is to divert someone's eyes away from whatever is it is they are doing at that time and towards your advert instead. You are competing with hundreds of miniature friend-faces on Facebook, clickbait listicles published by Buzzfeed, celebrity gossip news, and whatever else the distraction machine that is the internet has on offer that day. Therefore, you need to up your game and come in with the most attention-grabbing image you can conjure up—regardless of whether that image is related to your product. Trust me on this one—no one is going to come knocking on your door to complain that your advert image was peculiar. The most important thing is to get high advert click-through rates. Obviously, it's great if you find an attention-grabbing image that ties in with your product, but don't sweat it if you can't. You can always stamp your oddball image with a logo that communicates what you do through brain-dead obvious symbolism (e.g., a tiny picture of a bike says your company does something with bicycles).

Once you've grabbed a reader's attention with your interesting images, that reader will next scan your advert's headline and body. It is here that you clarify in words *what it is* you are selling.

www.dbooks.org

Headline

The purpose of the headline in advertisements is to indicate that this particular ad is for the special attention of the reader. In other words, your job is to communicate relevance. Imagine flipping through a newspaper. Most of the articles therein are of no interest to you. You are able to make that determination in split seconds. However, assuming that you are a web entrepreneur like me, as soon as you see the word "internet" in a headline or in some bolded text, you slow down and read in detail. (Of course, images can communicate this relevance too: I know I would slow down if I saw an artist's representation of the interconnecting nodes of the World Wide Web.)

The best headlines signal an advert's benefit to a *rushed* audience. The headline "1.5% financing interest rate" is the perfect fit for getting the attention of someone in the market for a mortgage—especially if your rate is lower than that of your competitors.

"Attn: Oxford Politics Students" is a great headline when advertising to students studying the same at Oxford. Indeed, this I'm-talking-to-you-style headline is effective *even if* your targeting settings limit your ads to *exclusively appear* to students studying politics at Oxford. This student audience has no way of knowing who you are targeting, and if you had used a more generic headline (e.g., "Attn: Politics Students"), the audience at Oxford might be inclined to think that your product is designed for a more general politics student audience (e.g., all politics students *around the world*), making your product appear like less of a glove-fit for Oxford students' needs, and therefore less attractive and interesting to them.

Body

The body is the persuasive part of the advertisement. Here you'll employ the regular principles of general copywriting and persuasion given in the previous chapter, albeit in severely character-limited form. If appropriate to your audience, consider sacrific-

ing grammar and spelling so as to cram as much sell per character as you humanly can. Txt spk is OK—at least sometimes.

Given that space is so limited, you might not have sufficient space to convince readers that they should buy your product. Here, your goal becomes convincing the reader to click through to your website and learn more. Of use here are tactics like arousing curiosity, inciting fear, imposing time limits on your offer, displaying social proof, or just knowing which of your product's benefits is the most compelling to that audience.

www.dbooks.org

21 | Paid Advertising: Strategic Principles

Topical Advertising vs Interruption Advertising

This distinction is about the state of mind of the person seeing the advert—regardless of how they were targeted.

The quintessential example of topical advertising is a Google Search ad which gets triggered based on certain search queries (e.g., "nearby gas station"). These ads reach people that know they have a *specific* problem and are in the market to solve it *right now*. The advert is fully on topic, so advertisers can expect high click-through rates, high conversion rates...and high competition from all the other businesses that have arrived at the same conclusion.

Interruption advertising, by contrast, is displayed when the recipient is focused on something unrelated to the subject of the advertisement. They are watching YouTube, reading an online newspaper, or dotingly browsing the Facebook profile of a secret crush. These adverts shoehorn their way into other experiences and demand your attention. Examples include ads between videos, flashy banner ads, or sponsored stories showing up in your Facebook feed. Interruption advertisements are the online marketing equivalent of cold-calling; the marketers deploying them have a difficult job, as can be seen by the resulting click-through rates that are often an order of magnitude lower than those achieved with topical advertisements.

I would recommend you pick topical advertising for your initial campaigns because it's conceptually simpler and usually more profitable (despite the high competition). But do not take this recommendation to mean that topical advertising is unequivocally better. There are cases when you're better off with interrup-

tion advertising. One is where your business sells a new product that is unknown to your market; because there's no way that anyone could be expected to specifically seek you out, you have no choice but to rely on interruption advertising, at least to begin with.

Another use-case for interruption advertising is when you've saturated your reach with topical advertising and now wish to extend your marketing to a larger audience only accessible through interruption advertising. Just because I don't explicitly google "Radiohead concert tickets" regularly doesn't mean that I wouldn't buy them if presented an opportunity. Interruption advertising ensures that this opportunity comes my way.

Rationed Views

Today there are over 7 billion people in the world. But unless you're selling water, only a small proportion of these people will have any interest in whatever it is you sell. This tiny group of potential customers is your intended audience, and it's these people that you target.

To a greater or lesser degree, your advertising only enables you to reach a narrow sliver of this targeted audience on any particular day. Imagine you are selling guides to pass law finals. Your targeted audience will be all students sitting their law finals at the end of this academic year. But only a fraction of these students will google "XYZ exam" on any given day, meaning that only a fraction of your audience will see the adverts you configured to appear for those keywords. Despite your targeting, your access to the underlying intended audience is rationed.

By contrast, suppose you had a parallel Facebook campaign that displayed adverts to Facebook users who identified themselves as studying law this year. Because a much larger percentage of law students log in to Facebook every day than search on Google for a particular law exam, you will reach a much wider slice of the underlying audience on any particular day. Through this targeting mechanism, your access is rationed to a much lesser degree.

www.dbooks.org

(Of course, you're unlikely to reach *everyone* in your targeted audience because some of them might not log in to Facebook *at all* while your campaigns are running.)

As with any attempt at binary categorisation, the full picture is messy; no targeting mechanism can be said to be fully rationed or completely unrationed. Day to day, membership of the Facebook law student group is mostly stable, with little in the way of migration in or out. If you were advertising to this group, nearly every member of the underlying audience would see your adverts after a week or so. By the end of that academic year, this audience would be totally saturated with your constant advertising…until, suddenly, the class disbands and gets renewed with the next batch of students, meaning that your advertisements become fresh again, at least for a week. As such, there are really two different rates of rationing, depending on whether you set your temporal resolution to days or years.

Despite its imperfections, the idea of rationing is a powerful one, and worth keeping in mind when thinking about your online advertising campaigns. Here are a few practical consequences and applications:

1. Possibility of gas pedals

Advertising to an unrationed group means that you can get through to the entire underlying audience in a short amount of time, so long as you're willing to bump up your budget sufficiently. Putting your foot firmly down on the gas pedal like this is great for launching a product or triggering a rapid-action bulge in sales. But, once everyone in your targeted audience has seen your adverts 10 times, your sales will peter out and your advertising campaign will cease to be effective. At this point, the revenue potential of this audience has been exhausted, at least until the underlying group refreshes its members. As an advertiser, you should anticipate this dip and be prepared to put a brake on your budget as soon as you notice the party ending, as indicated by plummeting click-through or conversion rates. In short, this is no fire-and-forget advertising strategy.

Contrast this strategy with the one applied to rationed groups. The advertiser here must accept that there is no gas pedal they can floor. Indeed, throttling their budget would be disastrous. In allocating a massive daily budget to their campaign, all they will succeed in doing is bombarding their thin daily ration of audience with the same ad hundreds of times. Whatever advertising budget isn't squandered in this way will remain unspent, sitting on the platform waiting for the next member of the target group to show up and be delivered the same advert for 110th time.

Advertising to a rationed group has advantages though, accruing to the patient advertiser. They will be rewarded with a yearlong trickle of cost-effective traffic. Because they can assume that the people experiencing advertisements on any given day are seeing them for the first time, they needn't worry about their adverts petering out due to saturation. Their campaigns will be low maintenance, and possible to leave unsupervised for long periods of time.

2. Targeting migration into the underlying group

The ideal targeting option focuses on *recent migrants* into your target group. The advantage is that these are the people that don't yet own your products or haven't yet seen your adverts. You'll be able to reduce advertising waste in proportion to your ability to specifically reach these incomers through specialised targeting.

Let's revisit our running example. If you had targeted *all* law students (from first year to finalist) instead of just finalists, then there would only be a partial refreshing of the underlying group each year, corresponding to the entrance of new first year students and exit of the finalists. But the old first year students would have become this year's second year students, and the old seconds this year's thirds, causing incomplete group renewal.

3. The need for anti-advertising-blindness measures

The human psyche is adept at tuning out irrelevant noise—literally, in the case of a rattly old air-conditioner whose din fades into the background noise; and figuratively, as happens when

www.dbooks.org

an advertisement ceases registering in our conscious awareness. This last phenomenon is known as "advertising blindness", and marketers who don't protect themselves from its effects will underperform compared to those who do.

Advertising blindness is more likely to turn up in unrationed advertising campaigns than their rationed brethren. When you squash so much advertisement into so little time, you increase the probability that the audience will be overexposed to the ads and will develop a perceptual blind spot towards them.

My own practical experience confirms the existence of advertising blindness. I noticed, when advertising to a relatively unrationed group on Facebook, that my campaign's click-through rates started to falter only three days after the campaign began. Luckily, this downward trend isn't inevitable—it can be reversed by introducing a few counter-measures.

The simplest of these is to create waves of advert variations. Once advertising blindness sets in, you disable the old adverts and swap them for the unseen variations. Since these new variants contain novel images and copy, they will be noticed afresh by your audience.

If you exit a room with a loud air-conditioner and return later that evening, you'll become aware of its noisy din once again. This is because your selective blindness resets after the stimulus is removed. This suggests that advertising blindness can also be lessened by turning adverts on or off, or by relying on frequency capping, the subject of our next section.

4. Frequency capping

Marketing lore has it that an ad must be seen *seven times* before it is fully effective. There are remarkably few studies done to confirm that seven is indeed the magic number, so take this as a matter of conjecture rather than fact. For example, Herbert E. Krugam found that maximum advertising response was reached

after just three viewings,[1] whereas Thomas Smith believes 20 are needed:[2]

> *The first time people look at any given ad, they don't even see it.*
> *The second time, they don't notice it.*
> *The third time, they are aware that it is there.*
> *The fourth time, they have a fleeting sense that they've seen it somewhere before.*
> *The fifth time, they actually read the ad.*
> *The sixth time, they thumb their nose at it.*
> *The seventh time, they start to get a little irritated with it.*
> *The eighth time, they start to think, "Here's that confounded ad again."*
> *The ninth time, they start to wonder if they're missing out on something.*
> *The tenth time, they ask their friends and neighbors if they've tried it.*
> *The eleventh time, they wonder how the company is paying for all these ads.*
> *The twelfth time, they start to think that it must be a good product.*
> *The thirteenth time, they start to feel the product has value.*
> *The fourteenth time, they start to remember wanting a product exactly like this for a long time.*
> *The fifteenth time, they start to yearn for it because they can't afford to buy it.*
> *The sixteenth time, they accept the fact that they will buy it sometime in the future.*
> *The seventeenth time, they make a note to buy the product.*
> *The eighteenth time, they curse their poverty for not allowing them to buy this terrific product.*
> *The nineteenth time, they count their money very carefully.*
> *The twentieth time prospects see the ad, they buy what is offering.*

[1] Herbert E. Krugman. "Why Three Exposures May Be Enough." Journal of Advertising Research 12, 6 (1972): 11-14

[2] Thomas Smith, *Successful Advertising*, 7th edn, 1885

www.dbooks.org

The ideal number of times an advert ought to be shown to a potential customer is more than once but certainly not "as many times as possible". There comes a point when you are just throwing money away. Moreover, aggressively frequent advertising irritates people, in much the same way as an overplayed TV rerun or Last Christmas on the radio.

Frequency capping, a feature available on some (but not all) platforms, limits how many times a given person will be shown your advert on any given day. Once that limit is reached, this person won't see your ads again, at least until tomorrow.

Frequency capping has an additional benefit beyond the ones mentioned already: It changes the *distribution* of your advertising reach. All things being equal, setting a low frequency cap (e.g., one impression per day) skews your budget towards showing your ads to as many different people as possible. Imagine the total size of your underlying target audience is 1,000,000 people but you only have the budget for 250,000 impressions over the lifetime of your advertising campaign. If you frequency cap at 1 impression per person, then your advertising platform can't rush through your budget by showing the same advert a thousand times to whoever happens to be online. Instead, the platform must wait until someone who hasn't yet seen your advert appears before it can display your advert and earn a commission.

When capping frequencies, keep in mind the *entire experience* of the person seeing the adverts. If you cap the frequency of one of your campaigns, this will not lessen the number of impressions from any parallel campaigns, either in the same or in other advertising platforms. Think about how your audiences do or don't overlap. Furthermore, even the lowest possible frequency cap, one impression per day, will oversaturate its audience if the campaign is left running for months on end.

While I am glad that frequency capping exists, I think advertising platforms can do a lot better. My beef is that they assume that the best time to show someone an advert again is...the next day! I beg to differ. Surely the ideal frequency would be the same as

the one trained mnemonists use for cementing facts into their long-term memories? If I learn a new word in French today, then memory experts advise me to revisit it tomorrow, next week, a month later, six month's later, and then two years after that. Why can't our advertising platforms follow the same tried-and-tested schedule?

The Progressively Rising Cost of Additional Leads

You are a retailer of pianos, the only one in your city. Today, 10 locals have taken the day off work and promised themselves to have ordered a piano by the day's end. On top of this, a further 100 people are *considering* buying one, but haven't made up their minds yet. They are still sitting on the fence.

The people who are determined to buy a piano *today* are likely to google with phrases that communicate immediate buying intention: "buy piano in Oxford", "price comparison Yamaha Clavinova". But those who are less decided will show less urgent and less direct interest through their search queries: "how to choose a piano", "which brand of piano is best for a beginner".

Advertising to those who want to buy today is a better use of your time and budget. These people don't need convincing—they just need a price and a confident promise of next-day delivery. All in all, it's an easier sales proposition. But there's not so many of them.

Once you have exhausted this initial supply of eager customers, revenue stops growing. Now it's up to you whether you want to branch out and advertise to the larger, second group of people who are vaguely considering buying a piano. With a sophisticated marketing campaign consisting not just of extended advertising campaigns, but also of weapons-grade sales copy and incentives like one-day-only discounts (that cut into your profit), you may convince one or two people on the fence to buy today. But all this effort has a cost, and the final tab shows that the price of selling to this second group is higher. The low-hanging fruit of the first group now exhausted, you must dip further and further into your

www.dbooks.org

pockets to continue making sales.

If the goal of business were to maximise margins and seek only the most *efficient* sales, there would be no point in financially straining to reach the customers in the second group. But this isn't the goal of business; the goal is to generate the most *overall profit*. As long as you continue to extract a healthy margin when taking into account the increased advertising costs, it's worth paying up.

This idea isn't always apparent at first sight because our emotions can lead us astray. To see what I mean, take a look at the results for these two advertising campaigns run on Google AdWords:

(Let's assume we are selling a product with no variable costs, such as an information product or software.)

Table 21.1 – "Google's Cut" is the perctange of your revenue that it charges you for advertising services

Campaign	Advertising Spend	Revenue	Google's Cut
Definites	£1,000	£3,000	33%
Maybes	£1,500	£2,000	75%

In the calculation above, you'll notice that advertising to the Definites is much cheaper. Google takes a 33% cut—a lot of money to be sure, but still a tolerable burden.

Now consider how you feel about the results of advertising to the Maybes. Your instinctive reflex might be to feel ripped off: "Why the hell should I give Google £1,500 for every £2,000 I earn? They're making more than me! I'm turning that rip-off Maybes campaign right off…. Good riddance!!" True to your word, you switch off the Maybe campaign and feel happy that those greedy a**holes at Google are no longer taking a bigger cut than you are.

A month later, you notice that your month's profits have sunk

20%. Dismayed, you analyse the situation once again, this time adding a "Profit" column to your spreadsheet.

Table 21.2 – "Profit" is simplistically defined as "Revenue - Advertising Spend"... remember that we have no variable costs here.

Campaign	Advertising Spend	Revenue	Google's Cut	Profit
Definites	£1,000	£3,000	33%	£2,000
Maybes	£1,500	£2,000	75%	£500
Total	£2,500	£5,000	50%	£2,500

As you can see, advertising to the Maybes generated a profit of £500.When this campaign was turned off, it wasn't like these leads would magically keep on visiting and buying. Without this campaign, the Maybes source of traffic froze to a halt, along with any profits they generated. By focusing on our outrage at Google's disproportionately large cut, we neglected our own profit figures. By letting our emotions take hold, we made a bad business decision. We're like the fool who, when offered "1% of £1,000,000,000" or "100% of £10", chose to take the 100%.

Let's continue this story. Now at peace with Google taking a large cut, your trust in the power of online advertising soars, so you decide to target a third group, the Question Marks. These are an even more speculative group than the Maybes, but you have reason to believe it's worth a shot.

After a month, the results are in, and they aren't good...

Campaign	Advertising Spend	Revenue	Google's Cut	Profit
Definites	£1,000	£3,000	33%	£2000
Maybes	£1,500	£2,000	75%	£500
Question Marks	£9,000	£6,000	150%	(£3,000)
Total	£11,500	£11,000	105%	(£500)

www.dbooks.org

Despite a large increase in revenue, profits have nosedived. Your business LOST money last month! Advertising to the Question Marks was a horribly failed experiment, so you react by radically decreasing the bids to this group, brainstorming alternative advert creative ideas, and maybe even deleting the whole campaign and just being done with it.

I depicted the effects of advertising to the Question Marks by neatly reporting them in a separate row. But this sort of orderliness isn't the reality for many advertisers. Through poorly implemented tracking technologies or lack of analysis know-how, these advertisers might not be able to distinctly group their customers into Definites, Maybes, and Question Marks in their performance readouts. All these advertisers have is an end-of-month summary of their Adwords performance (i.e., just the "Total" row from our table). And if that total is negative, as it would be this month, this advertiser may become an ex-advertiser, switching off Adwords for good.

This is a sad outcome, because if they had carved their audience into groups like we had, they would have seen that the Definites and Maybes generated a handy £2,500 profit per month—money they would not have otherwise earned. This shows us that decisions about advertising performance should not be made on a "this whole platform sucks" level. We should instead break up advertisements into specialised groups and reconfigure our campaigns to reinforce the successes and banish the failures.

Switch On All Bells and Whistles

Advertising platforms are always releasing special new features that change how your adverts will appear to leads. For instance, Google Adwords has what they call "ad extensions"[3]. These let you do things like insert 4–5 additional links into your ad; insert buttons for downloading an app; display a picture of a phone which mobile users can click on to call your office; add reviews; or insert "callouts", Google's term for short snippets of copy

[3]https://support.google.com/adwords/answer/2375499?hl=en

highlighting your key selling points (e.g., "free delivery", "100% money-back guarantee").

Unfortunately, there's no guarantee that your sitelinks will work after you enable them. Eligibility is dependent on your advertisement's rank and quality score, meaning that you may have some work to do before partaking in their advantages.[4]

Another way to stand out on Google Adwords is to choose an attention-grabbing ad format.[5] Instead of sticking with boring old text ads, you might try image, flash, gif, or video ads.

Twitter have their own parallel series of ad extensions via the same "Twitter Cards" we saw in the SEO section of this book. For example, the video card plants a video into a sponsored tweet, while the lead generation card inserts a call-to-action button that, when pressed, sends that lead's Twitter handle and email address to your company.

Turn as many of these extensions on as you can, if only for the primitive reason that they'll make your advert take up more space on the user's screen, thus grabbing you a larger share of their attention. As well as this "bigger is better" effect, extensions also give you another chance to convince your leads by means of providing social proof, extra links to your related products or services, playable videos, or simplified conversion funnels.

Unintentionally Competing with Yourself

Due to the auctioned nature of online advertising, it's possible for two of your adverts on the one platform to compete with each other and thereby drive up your advertising costs. The effect would be as if another business were bidding against you for that same audience.

By leveraging the organisational features of your online advertising platform, you should be able to avoid this problem. Ads

[4]http://services.google.com/fh/files/misc/settling-the-quality-score-whitepaper-final.pdf
[5]https://support.google.com/adwords/answer/1722124?hl=en

www.dbooks.org

within any "ad group" don't compete with one another, but do compete with ads from *other* ad groups. It is therefore advisable to avoid simultaneously targeting the same underlying audience through different ad sets.

Competing against yourself can happen in frustratingly subtle ways. For example, in Google Adwords, two keyword variations that seem distinct at first glance may in fact be fierce competitors. An ad group that targets broad-match keyword "sandwich grill" would compete with a separate ad group targeting the (also broad-match) keyword "panini grill", thanks to the synonym matching power of broad match.

Staff Working for Advertising Platforms Are Experts in Platform Features, Not in Your Profitability

(This is mostly a Google phenomenon, so what follows is Google-centric.)

Google isn't exactly known for its customer service. There is no way to contact its employees, come good or come bad. At best, you'll be invited to post your question to an official Google forum and pray that an empathetic onlooker happens upon your problem and advises you. But don't count on it.

Google's stronghold is impenetrable to mere mortals, except when those mere mortals splash out on buying Adwords. As soon as your budget reaches a certain threshold (approximately 500–1000 pounds/month), armies of keen Google representatives magically appear to offer you regular, free consultations on optimising your campaigns. Beginners will be walked through the most basic optimisations for quick wins, intermediate users will learn about recently released features that might nudge them ahead of their competition, and pros will be reminded to stop obsessing over the little things and refocus on getting the basics right.

I would do nothing but gush about my love for Google's free Ad-

words consultancy sessions, but for the problem of misaligned incentives. Google and its employees are paid whenever you spend money on advertising. The more you spend, the better it is for them. But from your point of view, the less you spend on advertising (for the same amount of sales), the lower your expenses and the higher your profits. These incentives are at odds with one another.

Now, don't get me wrong: Google seeks long-term win/win relationships with their clients, so they certainly desire that your advertising be *profitable* for you. But profitability is not an all-or-nothing matter: It's really a negotiation between you and Google about how much commission is reasonable in exchange for bringing in a sale. During these consultations, Google employees often advised me to raise my bids (i.e., to pay them more) in exchange for dubious "benefits", such as ranking above my competitors' ads or building brand awareness with hitherto unreachable customers, the ones for whom my previous low bids were insufficient to reach. But these advertising "benefits" were not worth the price they encouraged me to pay. And not once did Google advise me to *reduce* my bids—even when that was the wisest course of action.

www.dbooks.org

22 | Paid Advertising: Debugging and Optimisation

What Advertising Platforms Want

Online advertising space is sold by a rather strange species of auction. Whenever you bid for something in a typical auction—for example, with a real-estate auctioneer or on eBay—your bid is what you end up paying. With online advertising auctions, things are different: Your bid is the *maximum* you could be asked to pay, but your actual bill often turns out to be significantly less, at least when your particular mix of ad creatives, targeting, and ad relevancy helps an advertising platform maximise their total income.

Let's take Google AdWords as an example. Given that Google typically charges only after someone clicks on an advert, it stands to earn the most by showing adverts that have both high bids per click *and* high click-through rates. If I bid a million pounds per click for my ad, Google will eagerly display it to great swathes of their users, hoping for that million-pound pay-out. But if no one clicks my ad, Google eventually grows weary and thinks, "Screw this guy, we're not making any money on him because no one is ever going to click on his stupid ad". Google will gradually cease showing my ad and instead start showing someone else's, even if this other ad only offers to pay a paltry 70 pence per click. But if people actually click on this other ad (i.e., if it has a decent click-through rate), then a steady stream of 70-pence pay-outs will trickle into Google's coffers and it will end up earning more money in the end, despite the comparatively tiny bid offered for each click.

Many advertising platforms also care about maintaining a great user experience for their own users. Whenever a searcher thinks that Google's ads are useful and relevant, this increases the likeli-

hood that this searcher will continue choosing Google as their go-to search engine. For this reason, Google wants to encourage advertisers who create relevant adverts. This is done through its Quality Score metric. (Other platforms have equivalent concepts, such as Facebook's Relevance Score.) The higher the Quality Score an advertiser achieves, the greater their reach and the deeper the discount they earn in their advertising costs.

Stay Empirical

A word of warning before we get started: You should always check up on your campaigns after "optimising" them. Some time ago I studied all the best practices and integrated them into one of my campaigns. Having thought that I had brilliantly optimised my advertisements with all the latest tricks, I left the campaign alone for three months, feeling satisfied that I had done a great job.

But when I next logged in and looked over the results, I saw a 50% reduction in conversions. I learned an unsettling lesson that day: The best practices are not necessarily the best practices for *you*. In a field this complex, there aren't any certainties that you can apply without a second thought—and this is true even if the source of your best practices is the advertising platform itself (e.g., Google AdWords' official blog).

I should have been more sceptical and scheduled a check-up session soon after making my changes; I should have double-checked things with my own eyes. By not doing so earlier, I lost a lot of business.

How Much Should I Start Off Bidding With?

There is a dreadful uncertainty about how much to bid when beginning a new advertising campaign. With no historical data for you to refer to, you only have the platform's "recommended bid" to go by.

In business, as in life, only a fool pays the first price suggested to him in a negotiation. True to this truism, the initial bids recom-

www.dbooks.org

mended by online advertising platforms often inflate prices to a degree you would only expect at a bazaar situated in the most notorious of tourist traps. Knowing this, I recommend you set your opening bids to one quarter the amount the platform initially suggests. Because you are lowballing an algorithm rather than a human being, your aggressive negotiation tactics carry no risk of offending. You won't sabotage the deal, as might happen at the bazaar; the worst that can happen is your advertisements won't display until you raise your bids. Don't rush this: If your adverts are not displaying after an hour, that doesn't mean your bid is too low. Wait a day before making a better offer.

Why is it important to start low and raise your bids from there? Doesn't this have the same effect, you ask, as starting high then reducing your bids until they are just above the point where your adverts stop displaying? The disruptive factor which stops these two strategies from being equivalent is the presence of competition. Had you started with a high bid, your adverts would have displayed from the get-go, at the expense of your competitors' campaigns. These competitors would gradually notice that their advertisements had stopped showing, and then, intending to claw back their traffic, they would raise their bids above your initial, greatly inflated bids. Compare this to the situation where you start with a low bid and then raise your bids until they are marginally above your competitor's. Even after your competitor eventually responds to your stealing of their traffic—which might take weeks or even months—they will only nudge their bids slightly above yours. This slow tit-for-tat bidding war could go on for a long time, meaning that it may take months or even years before the bids are raised to the level they would have *started at* had you begun by bidding high (i.e., at the platform's "recommended bid"). By starting with low bids, you slow down competition-induced inflation of advertising costs. Or, to put in another way, whenever you start with a high bid, you prematurely raise the costs of advertising for everyone in that market.

What's the *Most* I Should Ever Bid?

As time goes on, bids for online advertising tend to rise and rise, mostly due to the unrelenting effect of mounting competition. There comes a point when you start asking yourself, "Should I raise my bids again, or have I gone far enough?"

The answer to this question is conceptually simple but technically difficult: Never bid more than your expected profit for that click, taking into account the minimum margin you wish to keep for your business efforts. This calculation requires that you have such excellent analytics and understanding of your business that you know the average profit you earn on each click. Without the ability to estimate this, you're flying blind.

Be warned that in certain crowded markets the most *you* can afford to bid will be insufficient to beat the bids of your competitors with deeper pockets. In this case, your only option is to become more efficient (e.g., by improving your advert click-through rates or on-site conversion rates, or by lowering your costs of production such that you can afford to pay more in advertising costs).

When Metrics Mislead You

Online advertising platforms provide you with a plethora of easily available metrics such as CPM (cost per 1,000 impressions), CPC (cost per click), and CPA (cost per acquisition). But when it comes to profitability, these metrics can be a bit loose with the truth. Unless you watch out for their pitfalls, they may lure you down expensive cul-de-sacs.

Without further ado, here they are, starting with the most distorting then ending with the least:

1. Misleading CTR (click-through rate)

The click-through rate measures the percentage of impressions that resulted in a click. It's a rough proxy for the compellingness of your advert vis-à-vis the audience seeing it.

www.dbooks.org

On its own, a high click-through rate isn't all that useful a metric. Its weakness is that it is totally divorced from your business's finances, making it pointless to optimise as an end in itself. If you have a mighty high CTR, say of 50%, this doesn't necessarily mean that your advert was a commercial success. The question of commercial success depends on what you *paid* for those clicks. If they cost you £100 a pop, then your advertising was a disaster, in spite of your sky-high CTR. If, on the other hand, the clicks only cost you £0.05 a pop, you're doing great. But in both cases, the metric that mattered was the CPC (or, if we're being more thorough, the cost per acquisition, etc.)

The main use of the CTR is in *comparing* the relative performance of your advert variants. This helps you determine which creatives were more compelling, arming you with the know-how to tweak and improve your other ads. These improvements in CTR should eventually lead to lower CPCs, thanks to mechanisms like Google's Quality Score and Facebook's Relevance Score.

2. Misleading CPM (cost per 1,000 impressions)

A cheap CPM is misleading when it's associated with a dismal click-through rate. Sure, your low CPM proves that you convinced the advertising platform to show lots and lots of people your advert for only a small sum of money. But if no one clicks on these ads, this also proves that these legions didn't care a damn for your adverts…those thousands of impressions left a dreadfully poor impression.

This combination of excellent CPM and terrible CTR could happen for all sorts of reasons: Perhaps your advert creatives were unconvincing. Perhaps your choice of targeting mechanism was *way* off and the wrong audience saw your adverts. Or perhaps your adverts were placed so far south in a scrollable page that no one even saw your advert, despite the advertising platform registering an "impression" and charging you for it.[1]

[1]Some advertising platforms have special features whereby they only charge you for an impression when an advert was *actually viewed*. For example, Google's Active View technology, released around 2014, only registers an impression when the advert was at least 50% visible

By the way, I'm not saying that having a cheap CPM combined with a low CTR is always bad. If the goal of your campaign is to establish your brand name in people's minds, this would be exactly what the doctor ordered.

3. Misleading CPC (cost per click)

First off, congratulations on creating a compelling advertising campaign that drives traffic to your website at a low cost! But before you start celebrating, take a quick look at your on-site conversion rates. If this low-cost CPC traffic isn't converting into paying customers at a healthy rate, you've got serious problems...

We need to separate out two possible pathologies before we go any further with our diagnosis. The first is the benign case where your on-site conversion rates suck *in general*. All your traffic, no matter what its source, has the same (dreadful) on-site conversion rate. In this case, your low-cost CPC isn't misleading at all. Although you have your work cut out for you in improving your overall on-site conversion rates, your cheap CPC has helped you make the most of your bad lot. You haven't been misled at all.

But if you discover that your low CPC traffic converts on-site at a worse rate than other traffic (e.g., traffic from other advertising campaigns or from SEO), then your attractively priced CPC has led you astray. This happens when there is a mismatch between what your advertisement promises (or suggests) and what you actually provide. If your adverts cause people to expect brand new laptops but you only sell used ones, the traffic you paid for will quickly abandon without converting. Likewise, if your audience expected free products but you only offer paid ones, they won't stick around for long.

These problems can be abated by clarifying your offer in your

on a user's screen for at least one second. Google also has a Display Network-only option (under "exclusions") which can exclude your adverts from displaying "below the fold". There's a theoretical argument that you should *always* use this exclusion, even when you are bidding on a CPC (cost per click) basis: Although there's no charge for unclicked CPC advertisements, having impressions appear below the fold decreases your advert's click-through rate, which in turn causes your Quality Score to suffer and your cost of advertising to rise.

www.dbooks.org

advert creatives. For example, if you're worried about too many freebie hunters clicking your adverts, you could include your product's price inside your advert body. This newly modified advert will almost certainly depress your CPC because now the people chasing freebies will be deterred by your price tag. But the clicks you *do* receive will be from people comfortable with paying for your products, and they are sure to convert at a higher rate. As your CPC drops, your CPA should soar.

In some cases, you'll need to change your targeting to repair mismatch problems. I lost more money that I'd like to admit by advertising my products in India and Pakistan. The problem was that my product wasn't suitable for these markets—they would have needed a local adaptation. I should have excluded these countries through targeting.

4. Misleading CPA (cost per acquisition)

Cost per acquisition represents the price you pay the advertising provider for each conversion (most typically, the price for each sale). The weakness of this metric is that it states nothing about the *magnitude* of the conversion. Maybe all those cheaply acquired conversions were for abnormally small basket sizes. You'll know this is happening if the basket sizes for this particular campaign compare unfavourably to your averages.

A friend of mine who owns an ecommerce store advertised a "loss leader" product. (This is when you sell a product at such a low price that you lose money; the goal is to convince people to come to your store and buy other, higher-margin products during the same visit). The advertising campaign for the loss leader had an excellent CPA, but after analysing the figures in detail he discovered that the customers won through this campaign purchased nothing other than the loss leader. The seductively low CPA figure obscured the commercial truth that his campaign only succeeded in reaching an unprofitable, bargain-hunter audience.

It must be emphasised that although a low CPA *can* mislead, often it won't. For example, businesses that have only a single possible

basket size (e.g., ebook sellers, retailers of paid apps) always have the same conversion magnitude, so therefore the CPA is a realistic and accurate measure of their advertising effectiveness.

5. Misleading ROAS (return on advertising spend)

ROAS is the amount of revenue generated by every pound spent in online advertising. This metric accounts for variations in basket size which are missed by the CPA, making ROAS a superior choice whenever you want to take the pulse of your advertising campaign's success.

But it is not perfect. One way an attractive (i.e., high) ROAS can mislead is when different orders have uneven *profitability* levels (e.g., because of the effect of variable tax rates, shipping costs, inventory costs, or some other factor). Suppose that Product A has a profit margin of 60% and Product B only 30%. An order for £100 of Product A would now be worth twice as much to your bottom line as an equivalently large order for Product B. But this difference would be completely missed by the ROAS, since this only sees the £100 in revenue earned in both cases.

There is a way to somewhat mitigate this shortcoming: Instead of basing the calculations on *revenue*, you base them on *profit*, or some approximation thereof. For example, you could subtract sales tax, shipping, and inventory costs from your gross revenue, then see how much of this reduced figure was generated by your advertising spend. This should give you a more accurate picture than basing the ROAS metric on simple revenue.

If you want to be finicky—which is generally a good thing when working with business metrics—you should know that even this improved ROAS has a blind spot: It fails to account for the total *lifetime value* of your visitors. Customers tend to become repeat customers, and any ROAS measurement linked to *individual* orders will fail to capture the *total* value a customer will bring to your business over the long term. Perhaps customers found through adverts on a website for a trade journal buy ten times as many widgets in the coming years compared to customers found

www.dbooks.org

through other ads. An ROAS measurement won't be able to tell you this, because its entire understanding is limited to the information contained within the initial order.

There are possibilities for integrating elements of lifetime value into your ROAS metrics, but these can be rather challenging to implement. One is configuring your website and Analytics to stitch together distinct user sessions belonging to the same person. Another is leveraging technologies like the Analytics Measurement Protocol to tie together future transactions and account for sales that happen outside the normal website channel. For more on these two, refer to the chapter on Google Analytics.

Overcoming Low Reach

Perhaps the most frustrating problem in the realm of online advertising is when your ads are unable to get enough airtime to spend your budget. Even the most brilliant advert ever concocted will fail to deliver any conversions if no one ever sees it. This failure to receive as many impressions as desired is known in advertiser jargon as having "low reach".

i. Targeting woes

When debugging low reach, the first thing to determine is where the fault lies. One possibility is that you targeted too narrow an audience and maxed out on all its advertising potential. The easiest solution here is to loosen your targeting—are you sure there aren't any unnecessary targeting restrictions in place? Once you've ruled this out, the next thing to do is seek out new targeting mechanisms in a bid to reach your intended audience from other, as of yet untried angles. You might target new keywords in Google Search, different subreddits over at Reddit, or just released real-world behaviours on Twitter. You might even try a new platform altogether, such as Snapchat Ads.

Roughly speaking, your search for new targeting angles will lead you to consider *more general* targeting options, *more specific* ones, and ones that are *laterally related* to your original idea. *"*More

general" would mean that instead of targeting "christmas cards", you target "gift cards" or even just "cards". "More specific" would mean the inverse; now you target "reindeer cards", "santa cards", and "christmas cards for kids". The goal is to accumulate a larger audience by targeting many specific niches. "Laterally related", then, just refers to any other relation that somehow connects back to your original targeting intention, even if that relation is tentative and tangential. Here you might target "christmas wrapping paper" in the speculative hope that there's enough overlap between wrapping paper and cards for this advertising to return a profit.

Because the number of things in the world is increasing—be they brand names, fields of study, Facebook pages, cultural memes, or newly minted words—there's always the opportunity to be the first mover into a targetable semantic space that other advertisers have yet to reach. By keeping up with the popular culture of your domain, you might be able to stay ahead of your competition.

ii. Overaggressive filters

Another targeting-related cause of low reach is the effect of inappropriately eager filters that act to categorically block your adverts from showing. I have in mind campaign-wide negative keywords, bans on adverts appearing on mobile devices, filters that prevent adverts from showing below the fold, and "exact match" keywords which are used when a broader match would have been more appropriate.

iii. Tough competition

Low reach can also occur when someone else is offering the advertising platform more money than you are to reach the same audience. You'll know this is the culprit when your metrics indicate that your targeted audience is plenty big, but you're only reaching a tiny proportion of it.

One such metric is Google AdWords' "impression share". This measures the proportion of times your adverts were shown relative to those of others competing for the same advertising slot.

www.dbooks.org

When your impression share is low, this tells you that your adverts aren't getting as much airtime as they could.

Over on Facebook you can infer that competitors are outbidding you by comparing the estimates of your reach (shown when keying in your targeting data) with the number of people you end up reaching. Large discrepancies between these two figures indicate that it's competitors and not your choice of targeting that is the cause of your problem. A word of warning though: You shouldn't be *too* eager to attribute your low reach to competitors. Because of the prevalence of ad blockers, a certain percentage of your audience—perhaps even the majority—may never see your ads, even after a decade of you trying to reach them.

Anyway, assuming that you're reasonably confident that your case of low reach is caused by competition, the solution is to increase your bids and optimise your advert creatives so as to improve click-through rates/Quality Score/ad relevancy.

iv. Low Quality Score

Google uses the Quality Score metric to decide whether your adverts are even *eligible* to appear for a given search query. As such, a low Quality Score can directly lead to a low reach. What's more, even if your Quality Score is high enough for your adverts to *appear*, it may not be sufficiently high for your ads to appear in any respectable position above the fold. This is because Quality Score also determines your ad's position. And because many web visitors never scroll down to the lower results in a Google search, this means that low-ranking advertisements will rarely be seen, thereby further reducing their reach. We'll discuss ways to improve Quality Score in a separate section below.

Other platforms have concepts similar to Quality Score (e.g., Facebook's Relevance Score), but it seems like these alternative metrics only affect your *cost* of advertising. There is no clear indicator that low quality scores on these platforms would prevent adverts from displaying in the same way that a low Quality Score on Google does.

v. Tightly rationed audience

This is the most benign cause of low reach. Sometimes your choice of targeting means that you'll only catch tiny glances of your underlying audience every day; it may be unrealistic to expect even half of them to appear online within the reach of your adverts after even a few months. The cure to this problem is straightforward: Simply leave your advertisements running for a longer time. This approach can be fine-tuned by taking steps to push your advertising platform into showing a greater proportion of your adverts to new people (e.g., by using built-in per-person frequency capping or creating negative remarketing lists populated with people who've already visited your website).

Those who study problem solving in the abstract have observed something that's relevant to everything we've said so far about increasing reach: *You can never really solve a problem; you can only replace it with another one—one that you (hopefully) consider more tolerable*. Whenever you take measures to increase your reach, you often decrease your advertisement's cost-effectiveness with that same stroke. You should anticipate needing to make this trade-off.

Overcoming Low CTR

Your click-through rate (CTR), calculated by dividing the number of times your advert was clicked by the number of times it was shown, is both a measure of the appropriateness of your targeting and of the compellingness of your advert creative in relation to that particular audience.

What counts as a good CTR depends entirely upon your business: 0.5% could be amazing for a mass-market ad for toothpaste, whereas 5% could be laughably bad in a niche market for yacht repair.

A better way to gauge success is to see how one particular ad's CTR compares to that of your other adverts. Suppose that you have eight adverts running in the one ad group, and that one

www.dbooks.org

of these ads has a CTR twice as high as any of the others. This indicates that there's something special about this one high-performing ad; it's your job to figure out what this special something is. Maybe you discover that the hyper-effective ad used a fear tactic in its headline, whereas all your other ads went for a "we're better value" tactic, which happened to be nowhere near as effective.

Another useful source of feedback is to measure your CTR against your competitors'. Google Display Network has a "Relative CTR" metric which tells you whether you perform better or worse than the other people competing for the same ad placements. A Relative CTR rate of x10 would mean that your CTR is ten times better than the average.[2]

The main commercial justification for improving CTRs is that it lowers the cost of advertising. Because of how ad auctions work, adverts with high CTRs pay less per click than those with lower rates.

i. Fine-tuning targeting

Perhaps the simplest way to improve your CTR is to tighten your targeting. If a proportion of the people you are advertising to have no interest in your product, then of course your click-through rate will suffer. The way out of this wastefulness is to study the results of your current advertising campaigns and identify any pockets of irrelevant targeting. For example, suppose you run a LASIK clinic (laser vision correction) and advertise with the broad match keyword "laser surgery". By studying your Google AdWords report of the *actual search queries* which triggered your ads, you would learn that your advertisements were showing on searches like "laser *tattoo* removal surgery" and "laser *wart* removal surgery". All the people searching for wart and tattoo removal have no chance of clicking on ads for laser eye surgery, so this would obviously drag down your advertising campaign's overall CTR. You could repair this problem by adding

[2]To view relative CTR in your AdWords account, go to the "Ad groups" tab, select "Customize columns" from the "Columns" drop-down menu, then select "Relative CTR."

negative keywords for "tattoo" and "wart", thereby surgically fine-tuning your targeting to exclude people who'll never click. A similar sort of optimisation is available on Google Display Network. Here, you would fine-tune your targeting by placing *negative* placements (i.e., telling Google not to show your adverts on websites or apps that previously gave you poor CTRs). Another option here is site/category exclusions (i.e., telling Google not to put your adverts on obituary websites or on parked domains, etc.) Yet another optimisation along these lines (and available in every platform) is geographic exclusion, whereby you stop advertising to people in regions with low CTRs.

ii. Completely different targeting

So far we've talked about fine-tuning your targeting with micro-adjustments. Sometimes this works, but other times it isn't enough to move the meter. In this case, it's possible that your current targeting mechanism is *way* off the mark and the only way to rescue your CTR is to try *drastically* different mechanisms (or to move to another advertising platform altogether). Instead of targeting by keywords, you might try targeting by workplace. Instead of targeting by workplace, you might try targeting by real-world behaviour. Instead of targeting on Twitter, you might try targeting on LinkedIn, etc.

iii. Form of advertisements

Once you believe that your targeting is in order, the next place to look for a CTR boost is within the *form* of your adverts. Are you using all the latest bells and whistles (e.g., AdWords Ad Extensions or Twitter Cards)? Are your ads viable to appear in the Facebook newsfeed with the maximum image size? Are your ads present on the latest platforms and devices?

Closely related to modifications of form are optimisations as to how your copy appears within adverts. For example, many advertisers Write Their Ads In Titlecase So As To Attract Attention. Or they put ^odd^ ***symbols inside their headlines and ad bodies. Or they stamp logos on all their images, thereby establishing brand trustworthiness.

www.dbooks.org

iv. Content of advert creatives

As soon as you've exhausted the possibilities for improving your CTR by modifying the *form* of your adverts, it's time to look at their *content*. This involves applying the techniques from the copywriting chapter so as to enhance your creative. A better image might help your ad stand out when it appears on a page crowded with competing advertisers; a catchier headline might convince more people to read on; and an alternative angle for your ad body might lure more people into clicking and visiting your website. Experimenting with advert creatives isn't particularly labour intensive, so there's no harm in trying out a range of ideas.

The ads of your established competitors form a particularly good source of inspiration. To take advantage of this treasure trove of ideas, load up Google's AdPreview tool, search for your keywords, then take notes on how competitors in other cities, countries, or niches are advertising.

v. Match between targeting and advert content

There is an interdependency between the *targeting* of your adverts and their *content*. It is only by matching the one with the other that you can achieve the highest possible click-through rates. For example, if your advert creatives emphasise the ability of your pocketknife to cut through fishing line, this might achieve an excellent CTR when targeting fishermen, but not so much when targeting wannabe tough guys.

The way to avoid these problems is to meticulously separate your adverts into small, thematic, tightly targeted ad groups. The idea behind this practice is that the advertising copy of each ad group should almost exactly match the sentiment and personality of the person likely to view the adverts contained within it.

These ad groups are smaller than you might think. Professional Google AdWords advertisers, such as Hanapin Marketing (PPC Hero), recommend having between 3–15 keywords in each

group.[3] Generally speaking, high value keywords (i.e., keywords that get a lot of impressions, clicks, or conversions) should appear in the smallest ad groups (i.e., of only 3–5 keywords). This helps you focus on optimising the heaviest hitters.

Overcoming Low Quality Score (Google-Specific)

Google issues your keywords a Quality Score on a 1–10 scale, with 1 being abysmal and 10 being God-like. It's worth improving your Quality Score because a good score leads to increased reach and decreased advertising costs.

According to Google, the big three components of your Quality Score are your expected CTR, your ad relevance, and your landing page experience. However, these factors are irrelevant when you launch new keywords which have only accrued relatively low numbers of impressions (i.e., less than a few thousand). Because Google has insufficient data from *your account* to give you a Quality Score, it averages the performance of *other advertisers* targeting those same keywords. This means that until you pass what's known as the "impression threshold", don't expect your Quality Score to respond to your efforts at improving it.

Assuming you've passed the impression threshold, what's the fastest way to improve your Quality Score? Since your expected CTR is the largest contributor to your Quality Score, the CTR is a good place to start. Google bases your expected CTR on your previous CTR performance for a given keyword (and for keywords related to it).[4] Therefore, your Quality Score will rise as a side-effect of applying the techniques we've already seen in the section about tackling a low CTR.

The next best way to improve your Quality Score is to work on your ad relevance. This is all about how closely your ad matches the intent behind a user's search. You'll get a lot of bang for your buck by studying the report about what actual search queries

[3] http://www.ppchero.com/16-commonly-asked-ppc-questions-and-answers/
[4] http://services.google.com/fh/files/misc/settling-the-quality-score-whitepaper-final.pdf

www.dbooks.org

people make and ensuring that these same queries appear somewhere in your advert creative (be that in its title or its copy). It's also important to ensure that you don't have too many disparate keywords per ad group. A wide range of keywords cannot possibly address any one user's query directly, and will be sure to drag down your ad relevance. Again, this advice is very similar to what we suggested for improving low CTRs above.

The last way to improve your Quality Score is by polishing your landing page experience. The important factors here are page load times, transparency and trustworthiness (e.g., giving your contact details on your website and explaining what you do with a user's data), navigability/usability (e.g., not having obnoxious pop-ups, and being mobile-friendly), and the presence of relevant content (i.e., your landing page should match the content of your advert and keywords).[5] An effective way to ensure relevancy is to create landing pages specifically for each ad group.

If you're unsure which of the big three components of Quality Score is most problematic in your specific case, Google gives guidance through a feedback tool integrated into the AdWords platform.[6]

Overcoming Low Relevance Score (Facebook-Specific)

Facebook calculates your advert's Relevance Score by comparing your advert's expected positive feedback against its expected negative feedback. Positive feedback includes how much Facebook expects people to like your advert, share it, view the contained video, or click on it. Negative feedback is how often Facebook expects people to hide your advert from their timeline or report your video.

A high Relevance Score significantly decreases your cost of advertising, so it's obviously worth improving it. One way is by

[5] https://support.google.com/adwords/answer/2404197?hl=en

[6] Log in to AdWords. Go to the Keywords tab of your campaign, then hover your mouse pointer over the "status" speech bubble for any keyword. A pop-up will appear detailing how you perform in each Quality Score dimension.

increasing your CTR, a topic we've repeatedly covered earlier. Another way is to follow Facebook's recommendations to 1) use laser-narrow targeting, 2) use their special call-to-action buttons, and 3) give details about your business within the advert creative.[7]

Yet another way to improve your Relevance Score is to refresh your advert creatives often; Facebook members will get annoyed if they see the same advert a hundred times. I know I'd consider hiding an ad at that point.

Perhaps the most interesting point of leverage you have in improving your Relevance Score is to take account of the social aspects of Facebook. Funny or quirky photos are going to get liked, while offensive or inappropriate photos are going to get hidden.

Overcoming High CPC

(I'll use the term "CPC" here to refer to either "cost per click" or "cost per 1,000" impressions [CPM], since the ideas are the same.)

A high CPC burns through your advertising budget faster, yet gives you less for it. You reach less people and trigger less sales than would have been possible if you had a lower CPC. Indeed, the ongoing viability of your paid advertising campaigns depends on your CPC staying lower than the average profit you earn per click.

The simplest way to reduce your CPC is to bid less. As we discussed in a previous chapter, there is a progressively rising cost of advertising to each additional person. By reducing your bids, you skip the most expensive-to-reach people in any given audience, and, as such, the high cost of advertising to these few won't spill over and inflate your overall cost of advertising.

The trade-off when bidding less is that you'll reach fewer people. For this exact reason, reducing your bids is sometimes a dreadful idea. This is so in cases where your revenue is limited by the number of customers, and where the prices you pay per click are

[7] https://www.facebook.com/business/a/ads-relevance-score

www.dbooks.org

still less than the profit you expect to earn for it. Here, you want to do everything in your power to gain as many customers as you can—including bidding more.

The other ways to reduce CPC are nothing new; they involve techniques we've already seen, such as experimenting with different targeting mechanisms, improving your CTR, and increasing your Quality Score.

Overcoming High CPA

Your cost per acquisition (CPA) depends on the entirety of your conversion funnel, which includes not just your advertisements themselves, but also your website's ability to convert this paid-for traffic into paying customers.

There's not much new to add at this point in the book. We've already looked at how to optimise *adverts* in this chapter, and we looked at techniques for improving *on-site conversion rates* in a previous one. The only remaining factor to discuss is the need for a harmonious *connection* between your adverts and your website.

Imagine you own a site that generally has a great on-site conversion rate. You decide to send it extra traffic through paid adverts which themselves have impressive CTRs and CPCs. You expect that this combination of effective adverts and an effective on-site conversion rate should culminate in an especially effective CPA. But this effectiveness fails to materialise; for some reason you end up with a dismal CPA. What went wrong?

After some analysis, you learn that your on-site conversion rates, while normally excellent, were appalling in this one case. This discrepancy tells you something important: Your adverts somehow didn't match your website. Perhaps your ads promised something that your website couldn't deliver. Perhaps they sent you traffic that speaks another language or cannot view your pages from their browsing devices. Or perhaps your advert suggested that your product was free when it wasn't. The possibilities for mismatch are endless, and the only way to improve your CPA in

this case is to take out your detective hat and start searching for clues.

The End

I'm flat-out flattered that you've read my book all the way to the end. I take the highest of compliments from your taking the time out of your busy life to listen to my thoughts for so long. I sincerely hope that you took something valuable away from this book that will help you thrive in your own projects, whatever they might be.

In a weird way, I feel that we know each other, asymmetric as that may be.

Review Request

Reviews and word of mouth are the lifeblood of any book's commercial success, so you'd be working a small miracle for me by rating this book on Amazon or GoodReads (click here), tweeting about it to your followers, or mentioning it whenever you see someone asking questions about online marketing.

www.dbooks.org