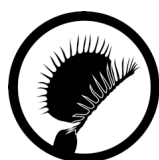


DECEPTIVE PATTERNS

EXPOSING THE TRICKS TECH COMPANIES USE
TO CONTROL YOU

FREE SAMPLE CHAPTERS

HARRY BRIGNULL



In memory of Mike Scaife, whose mentorship set me on this path.

ACKNOWLEDGMENTS

I am extremely thankful to the kind folk who reviewed early versions of the manuscript and helped me improve it.

Professor Yvonne Rogers introduced me to HCI and interaction design many years ago and taught me throughout my Undergraduate degree, Masters and PhD. She also helped me with this book. Thank you, Yvonne.

The legal response to deceptive patterns is rapidly evolving, especially in the European Union. While in no sense is this book a legal text, I would like to thank Dr Mark Leiser at Vrije Universiteit Amsterdam for helping me make some sense of the legal landscape for regulating deceptive design.

I'd like to thank Cennydd Bowles who spent hours helping me unpick and rewrite some of the more tricky arguments in the book. He also helped me with the first talk I ever wrote on this topic, more than a decade ago.

I'd also like to thank Finn Lützow-Holm who guided me towards lots of useful themes and research. It's become such a big area and Finn seems to be tapped into all of it.

Finally, I had a number of reviewers who asked to remain anonymous. You know who you are. Thank you.

Editor: Owen Gregory

Cover Art: Mia Brignull

ERRATA AND FEEDBACK

This is a first edition and no mistakes or omissions have been identified yet. If you wish to submit or review the known errata, please visit the webpage below. Similarly, if you work for an organisation that is mentioned in this book and you wish to provide a response, please use the same page. Thank you.

www.deceptive.design/book-errata

Copyright © 2023 by Dr. Harry Brignull

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.

The information provided in this book is intended for general informational purposes only. While certain legal topics are discussed, it is crucial to note that this book does not constitute legal advice. The content presented herein is not a substitute for professional legal counsel, and readers are advised to consult with qualified legal professionals for specific legal concerns or queries. The author and the publishing company disclaims any liability or responsibility for any actions taken based on the information provided in this book. Readers are encouraged to use their own discretion and judgement when applying the concepts and principles discussed herein.

Published by Testimonium Ltd

First Edition: 2023

PROLOGUE

Their faces give nothing away. It's a Thursday afternoon in March 2021, and the Communications and Technology subcommittee of the 117th Congress is holding a joint hearing online. Three of the world's most powerful people have been invited to give testimony in a session called 'Disinformation Nation: Social Media's Role in Promoting Extremism and Misinformation'.¹ Sundar Pichai, CEO of Google; Jack Dorsey, CEO of Twitter; and Mark Zuckerberg, Chairman and CEO of Facebook.

It's the moment I've been waiting for. The camera cuts to House Representative Lisa Blunt Rochester. She introduces the concept of dark patterns and defines them as 'intentionally deceptive user interfaces that trick people'. She asks Pichai, Dorsey and Zuckerberg:

'Would you oppose legislation that bans the use of intentionally manipulative design techniques that trick users into giving up their personal information?'

As the camera cuts to each of the CEOs, we see a stark difference. Lisa Blunt Rochester is sitting in a tiny wooden booth, connected with a grainy laptop webcam, but each one of the CEOs is evidently on a film set with professional lighting, cameras and microphones.

2 DECEPTIVE PATTERNS

Picahi replies promisingly, 'We definitely are happy to have oversight on these areas.'

Dorsey replies with just three words 'Open to it.'

Zukerberg is more evasive. 'Congresswoman, I think the principle makes sense and the details matter.'

His reply seems to antagonise Blunt Rochester, who pushes him further: 'OK. Mr Zuckerberg, your company recently conducted this massive ad campaign on how far the internet has come in the last 25 years. Great ad. You ended with a statement: "We support updated internet regulations to address today's challenges." Unfortunately, the proposal that you direct your viewers to fails to address dark patterns, user manipulation, or deceptive design choices. Mr Zuckerberg, will you commit now to include deceptive design choices as part of your platform for better internet regulations?'

Zuckerberg hesitates for a moment: 'Congresswoman I'll... I'll think about it. My initial response is that I feel there are other areas that I think might be more urgently in need...'

Blunt Rochester cuts him off and gives a final speech, knowing her five minutes are almost up. 'If you say this is a desire of yours to address the issues that we face today – dark patterns goes back to 2010 – this whole issue of deceptive practices. And I hope that you will look into it! I will say [...] our children [...] our seniors, veterans, people of color, even our very democracy is at stake here. We must act and I assure you – we will assure you – we will act.'

A moving speech, but the CEOs are holding all their cards close to their chest. They know regulatory change is coming, but they don't want to give away any more than they have to.

Lisa Blunt Rochester was spot on in her statement. The concept of dark patterns harks back to early 2010. I know this, because I coined the term; though had I known it would become so popular, I would have

taken a bit more care with the name. I remember sitting at my kitchen table in May 2010, ballpoint pen in hand. As I wrote about this topic for the first time, I was putting together a talk for a conference. *'I'm not sure there'll be enough here for a 20-minute presentation,'* I thought to myself – but the more I looked, the more I found. Deceptive tricks and techniques were in use all over the place and, at the time, nobody was talking about them.

A lot has changed since then.

PART ONE
DIVING INTO THE WORLD OF
DECEPTION

CHAPTER 1

INTRODUCTION

In 2010, I defined a *dark pattern* as: ‘a user interface that has been carefully crafted to trick users into doing things, such as buying insurance with their purchase or signing up for recurring bills’.

This definition is now a little out of date, and today I prefer to use the term *deceptive pattern*,¹ or to be pedantic, *deceptive or manipulative pattern* – but that’s a bit of a mouthful, so in this book I’ll use *deceptive pattern* as a shorthand to mean both.²

At the time, I was probably the only researcher looking closely at the area of manipulative and deceptive user interface design. Now, over thirteen years later, the area has blossomed into a multidisciplinary topic involving numerous human–computer interaction (HCI) researchers, legal scholars and many other people. Of course, I can’t take credit for the work they’ve done; although I launched the initiative and defined a dozen or so of the initial terms, my role since then has mainly been that of an educator, campaigner and amplifier³. I’ve worked to spread awareness, to name and shame companies, and to encourage legislators, regulators and enforcers to take action.

To understand how businesses can employ design to manipulate users for profit, let’s start with a physical example: travelling through an

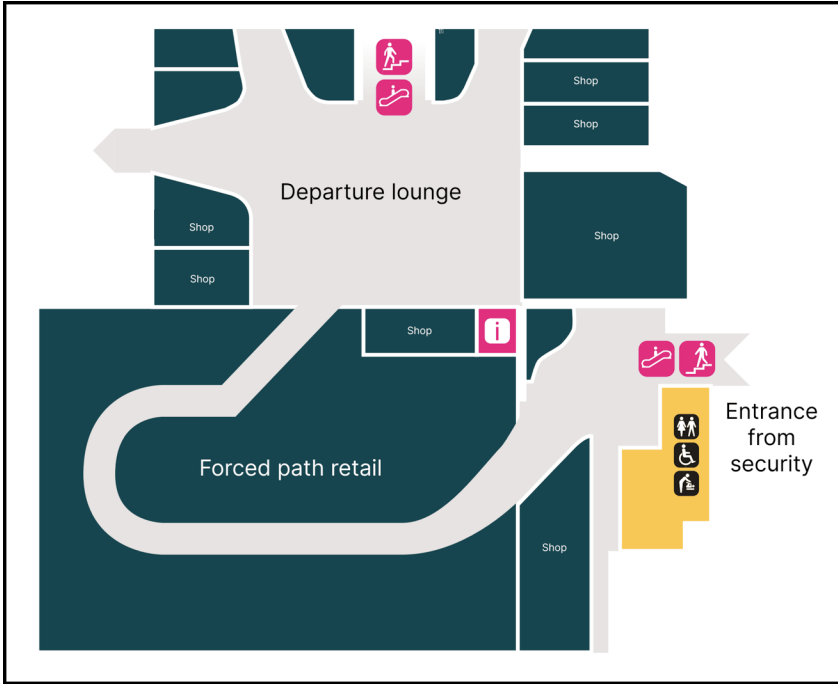
8 DECEPTIVE PATTERNS

airport. When you travel through London Gatwick Airport, you're advised to 'arrive at least two hours before your flight to allow plenty of extra time to check-in and pass through security.'⁴ But after you go through security at Gatwick, you're not allowed to go directly to the departure lounge. You're forced to do something that has nothing to do with your trip, and it consumes your attention, energy and time. You have no choice in the matter – even if you're running late.



The London Gatwick mandatory retail experience.

In the industry, this is known as a 'forced path' store layout.⁵ It's really just a shop that's a long, winding corridor, packed into a rectangular footprint in the same way your gut is packed into your belly – travellers are forced in one end and come out the other. The curved path serves a useful function for the business – it forces retail displays into the centre of the traveller's vision, making it almost impossible for them to avoid looking at the stuff on sale as they navigate their way through the area.⁶



Floor plan of the London Gatwick south terminal, featuring a mandatory forced path that doubles-back on itself.

Think for a moment about the airline tickets and legal terms. In those documents, there's nothing mentioned about requiring you to spend time in a retail area looking at perfumes, beauty products and alcohol before you're allowed into the departure lounge. And consider the airport's guidance – to arrive at least two hours before your flight. If time efficiency really was their top priority, they wouldn't impose the forced path retail store as a mandatory step between security and the departure lounge.

This is a good example of how businesses can use design to coerce and manipulate you. Arguably, it's also slightly deceptive in the way that the business is fully aware of the revenue-generating purpose of the forced path store, yet they don't mention it when they ask you to arrive two hours early, and they don't give you a shortcut to skip it.

In this example, the negative impact on travellers is minor and not particularly harmful; it's more of a nuisance than anything else. But when you consider the fact that over 40 million people travel through Gatwick every year, you can see why it's designed this way.⁷ If this manipulative design can get just a few percent of travellers to make a purchase who would not otherwise have done, the airport can charge a huge premium on the lease for that retail space and enjoy a lucrative relationship with the retailer.

It's even easier to build manipulative and deceptive experiences online, because the designer has so much more within their control. When everything is virtual, anything can be tweaked to increase profitability. Here's a simple example of a deceptive pattern on a website. You've probably run into something like this yourself before when signing up to something:⁸

WIRED and Conde Nast would like to contact you with offers and opportunities. Please tick here if you would prefer to receive these messages:
by email ☐ by SMS ☐

If you do not want to hear from us about other relevant offers, please tick here:
by post ☐ by phone ☐

Our partners sometimes have special offers that we think you will find relevant, please tick here if you would prefer to receive these messages:
by email ☐ by SMS ☐

Please tick here if you would prefer not to hear from our partners:
by post ☐ by phone ☐

Excerpt from the Condé Nast Wired Magazine sign-up form (October 2010).

Did you see the trick? There's a switch in the wording between each line of checkboxes. If you tick the boxes in the first row, you're opting in to messages. In the second row, you tick them to opt out. Third row is opt in again, and fourth row is opt out. If you want to opt out but you're not paying attention, chances are you'll misunderstand at least

one of the rows and end up getting spammed. This trick enabled Condé Nast to send out more marketing messages, which meant more ‘eyeballs’ – more people seeing the information – which in turn meant more sales and more profit. If you live in the EU or the UK, you probably haven’t seen this type of deceptive pattern recently because it became illegal under the General Data Protection Regulation (GDPR)⁹ a few years ago.¹⁰ Hooray for progress!

Part of the inspiration for my work on deceptive patterns came from an interest in *design patterns*. A design pattern is a common and reusable solution for a problem when you’re building user interfaces (UIs). For example, if I told you to close your eyes and imagine the sign-in box for a website, you’d probably see the same thing in your mind’s eye as I do – a text field where you’d type your username, a password field below it, some kind of button that says ‘sign in’ and a link that says ‘Forgotten password?’. That’s a UI design pattern. Different industries have their own design patterns, and the idea originally comes from architecture in the built environment.¹¹

Another well-known idea is the *antipattern*: a common mistake when trying to solve a problem. But as I sat there, back in 2010, doodling in the margins, I realised there was another type of design pattern that nobody was talking about. It wasn’t about recommended practices or mistakes to avoid – it was about manipulative or deceptive practices that benefit the businesses that employed them and harmed the users who fell victim to them.

Although it’s taken a long time, this area of work is finally seeing the fruit of its labour as new laws emerge. We now have the EU GDPR, Unfair Commercial Practices Directive (UCPD), Digital Markets Act (DMA),¹² Digital Services Act (DSA),¹³ the proposed EU Data Act,¹⁴ the California Privacy Rights Act (CPRA),¹⁵ and the Colorado Privacy Act (CPA).¹⁶

The CPRA and CPA both use the same definition: ‘dark pattern means a user interface designed or manipulated with the substantial effect of subverting or impairing user autonomy, decision making, or choice’.

Central to this definition is the concept of autonomy – for a user to be able to act according to their own goals, free from external influences, while understanding the nature of their choices. For example, if a user is tricked into sharing personal information because the legal agreement was completely hidden from them, then by definition there is no agreement: the user was denied their autonomy, since they were not free to become informed and make their own choices. However, the CPRA and CPA only cover privacy. The United States doesn't yet have any state or federal laws that directly address deceptive patterns beyond privacy. The EU is slightly ahead in this regard, with the much broader Digital Markets Act and Digital Services Act coming into force in 2023. The DSA uses the following definition (Recital 67):

'Dark patterns on online interfaces of online platforms are practices that materially distort or impair, either on purpose or in effect, the ability of recipients of the service to make autonomous and informed choices or decisions. Those practices can be used to persuade the recipients of the service to engage in unwanted behaviours or into undesired decisions which have negative consequences for them.'

As you can see, the DSA's definition is similar to the CPRA and CPA. It's about not interfering with users' autonomy, choice and decision-making.

There are a few different ways to think about deceptive patterns, and the legal perspective is just one of them. For example, if your background is UI design or engineering, you may be more interested in the mechanics of how they're put together. If you're coming from psychology or HCI then you may be more interested in how they prey on the human mind. If you're an ethicist then you may be interested in the broader philosophical implications. In the coming chapters, this book will touch on each of these perspectives.

My main point here is that deceptive patterns are not just a niche curiosity any more. If you work in the tech industry you need to understand them, particularly since some types are already illegal, with even more activity coming from lawmakers, regulators and enforcers.¹⁷

Before we go much further, you'll need an understanding of how the design industry has evolved, too.

CHAPTER 2

A PRIMER ON DESIGN INDUSTRY TERMINOLOGY

It's easy to think of design as how things look. Fonts, colours, textures, grids, mood boards – that sort of thing. This is *graphic design*: it's still important in its own way, but it's now just a small part of what the digital design industry has become.

Today, design is far less about how you decorate things, and far more about how you persuade and influence people into doing things. It's mainly about tracking, testing, psychology, behavioural economics, statistics and empirical scientific research. In other words, it's all about achieving business goals and making money.

You might not realise it, but when you use popular apps or websites, the details of everything you click on and scroll through usually gets recorded. Then it gets analysed, carefully. In big companies like Meta, Amazon, Netflix and Google, they have teams of people paid six-figure salaries, tasked to work out how to make more money out of you. Every day, your behaviour is tracked and you take part in quantitative research (e.g. 'A/B tests' or 'multivariate tests') to work out what will make you click, buy or agree to the legal terms. It's important to understand that the same research methodologies can be used to help or harm users. It depends on the intent of the business owner. It just so happens that deceptive patterns are easy to build and deliver measur-

able outcomes, so deception is commonplace unless a business owner takes a strong position on preventing it from happening.

Deceptive patterns aren't always the result of rigorous research and careful craftsmanship – sometimes they're just profitable accidents. Consider the example of a subscription offer that doesn't clearly explain the nature of the ongoing charges, just because the writer didn't take due care. This might result in a surge of revenue, which the business may then come to rely on, and they may not even understand why.

I'm going to use some industry terms in this book, so I'll define them here.

PRODUCT

This is the general term that's used to refer to an app or a website or any other piece of software that people use. The Amazon app is a product. So is the Facebook website. You get the general idea. Sometimes companies prefer to refer to their business as offering a 'service', particularly if it involves customers interacting with different people and numerous touchpoints over a period of time.

PRODUCT MANAGERS

In most modern organisations, a single individual is directly responsible for all of the decision-making for a given product or feature. This person is known as the product manager (PM). They're usually like a mini CEO, responsible for everything within the realm assigned to them, though the exact title and job description varies. If a deceptive pattern is created, then the PM of that product should know about it. They should know why it's been created, what purpose it serves, how many users interact with it and how it makes money. This is handy to know if you're ever involved in choosing who to subpoena in a class action lawsuit.

USERS

A user is the category of person for whom the product is intended, rather than ‘all humans on the planet’. In the industry, we sometimes say *active users* for people who regularly use a product, and *target users* to include those for whom it is intended, but who might not be using it yet. The terms ‘monthly active users’ (MAU) and ‘daily active users’ (DAU) are also commonly used when measuring the success of a product, and deceptive patterns are often used to boost these numbers.

USER INTERFACE DESIGN

An interface is the point at which two things meet and interact. If you glue two pieces of wood together, the glue is the interface. In this case, instead of having two pieces of wood, you have a product and a user. The glue in the middle is the user interface (UI). With a screen-based device, we’re mainly talking about text, images, boxes and buttons: these components make up the user interface. With a voice-oriented device, like an Amazon Echo, the user interface is the words or audio that comes out of its speaker, and the commands it recognises when you speak into its microphone.

USER EXPERIENCE DESIGN

A user experience (UX) is what you perceive or feel when you interact with a product’s user interface over a period of time. If the interface is hard to use, then you’ll have a negative experience.

However, not all user experiences have the same strategic goals. For example, when you pay for something online, you want the checkout to be pain-free and quick. Most form-filling experiences are like this – you don’t want it to be fun, you want it to be done. In this context, usability and efficiency are paramount. Conversely, when you switch on a Nintendo or put on an Oculus headset, you want to savour every moment of the experience. In this context, emotions and entertainment matter.

Of course, there are many other kinds of human endeavour that need different design considerations. If you're designing an educational product, you need to understand how people learn. If you're designing the controls for an X-ray machine, safety is one of your biggest concerns. The list goes on and on. It's the job of a UX designer to think about these things. A UX designer takes a business's goals and marries them up with an understanding of user needs and user psychology. UX designers typically create sketches, diagrams and models – things that help with thinking and collaboration, forming a bridge between the people in the different roles in their team: product managers, researchers, technical subject matter experts and UI designers.

Unfortunately, the design industry has very few universally recognised certifications, or universally defined job titles, roles and responsibilities. Each company tends to use slightly different terminology and processes.

ALTERNATIVE TERMS FOR DECEPTIVE PATTERNS

Although the term *dark pattern* is still in use by some people, we should aim to phase it out and use more inclusive terminology that avoids negative associations. My preferred term is *deceptive pattern*, although if I am working with lawyers, I use the longer term *deceptive or manipulative pattern*, since not all of these patterns are deceptive. Various groups around the world use different terms to mean broadly the same thing:

- **harmful online choice architecture:** this term is used by the UK's Competitions and Markets Authority (CMA).
- **asshole design:** a colloquial term, used on Reddit and other forums.
- **dark nudge:** this term is sometimes used by behavioural economists, building on Richard Thaler and Cass Sustein's term 'nudge'.
- **sludge:** a term that specifically refers to obstructive design, which Cass Sustein has written about extensively.

It is unlikely we'll reach a universally agreed term any time soon, since this area of work now overlaps with legislation and regulation. For example, the word *deceptive* has a narrow technical definition in the United States at a federal level (due to the FTC Act), so the term *deceptive pattern* would be used very cautiously by US legal professionals (unlike in this book where I use it as a broad term).¹ Similarly, *dark pattern* has recently been defined in EU law, so it will continue to be used there despite its shortcomings. My view is that if you're not a lawyer or involved in legal systems, it's sensible to just be clear and descriptive about the design patterns you are talking about, and accept that there may be some movement in the terminology for this stuff as time passes.

CHAPTER 3

THE RISE OF DECEPTIVE PATTERNS

When I started working on deceptive patterns, I was a little naive. I thought they might be eradicated if we could name and shame the companies that use these practices. Or at the very least, perhaps we could encourage UI and UX designers to use a code of ethics that would reduce the number of deceptive patterns in existence.

This approach didn't work. In fact, things have become a lot worse since then. Deceptive patterns are everywhere now – there's even a tip line that takes reports from concerned users and relays them to policy-makers and enforcers around the world.¹ The fact we need a tip line at all means there's clearly more to do.

To be fair though, deceptive patterns didn't appear overnight. Deception is part of being human – in fact, it's so common in the animal kingdom that we even can think of deception as a feature of life itself.² The cover of this book features a Venus flytrap (*Dionaea muscipula*). This plant releases a scent that mimics the bouquet of fruits and flowers. Insects are attracted, and when they touch its sensory hairs inside the jaws, it snaps shut and traps the prey. This image is intended to be emblematic of unscrupulous tech companies who trick and trap their users using deceptive patterns.

Many historical stories and myths revolve around deception, such as ‘taking the King’s shilling’. In the 18th and 19th centuries, Britain spent a lot of time at war. But a career in the army or navy during wartime was not very attractive. With volunteers short on the ground, press gangs emerged to aggressively encourage recruitment, offering a shilling for every man who joined up. As the story goes, the act of receiving the coin was seen as a binding agreement, so unscrupulous recruiters would slip the coin into a sailor’s pocket or tankard of beer. When it entered their possession, the deal was done, and the men would be forced into naval servitude. Myth or not, the analogy with deceptive patterns is a strong one. Whether it’s clicking an ambiguously labelled button in a user interface or receiving a drink containing a hidden coin, it’s obvious that there’s a problem with the definition of consent if a person has no recourse after such a small, unintentional act.

It’s useful to think about what makes commercial deception and manipulation different today versus the pre-internet era. There are some aspects of modern technology that have acted as an accelerant or a catalyst, intensifying and spreading these practices.

THE RISE OF METRICS-DRIVEN CULTURE

The idea of being driven by metrics dates back a long way: there’s archeological evidence of accounting records from Mesopotamia, 7,000 years ago. Crude as it may have been then, human beings have got better at measuring things over time, and we’re now fanatical about measuring things accurately.

What’s changed is that the barrier to measuring things is now much, much lower. You don’t need to be particularly clever or have a lot of capital to start measuring anything and everything you do in a business environment, and to start using data analysis to inform your business decisions.

In fact, metrics-driven management can be quite easy. You work out what metrics matter to your business, then you reward your teams for

pursuing them using management techniques like performance-related pay, target metrics, bonuses and promotions. Of course, rewarding people for meeting a goal is almost the same as punishing them for not meeting it. In countries with less stringent labour standards, some companies use a management technique called ‘stack ranking’. This involves rating employees according to their performance on various measures, arranging them in rank order and then getting rid of the lowest performers. If an employee’s healthcare or immigration status is tied to their continued employment, this creates an enormous pressure on employees to do anything they can to hit their targets.

The web has also made it much easier to build and optimise deceptive patterns. With that in mind, I’d attribute the rise of deceptive patterns in software to the following general factors.

EASIER TRACKING

Before the internet, it wasn’t easy to observe people without them being painfully aware of being watched. The traditional observation method was to send researchers to a store and have them stand there with a clipboard.³

But field researchers are costly and can only look at one thing at a time. Today, all you need to do is add a snippet of JavaScript to your website to get in-depth tracking that observes every conceivable behaviour of every user of your product simultaneously, and have it recorded into a huge database in the cloud. Business owners have also noticed another advantage to online tracking. Despite it being more invasive than ever before, people don’t feel anywhere near as worried about their privacy being invaded – because they don’t feel human eyes on them. All that tracking happens behind the scenes, out of sight and out of mind.

Then you’ve got the data processing. Before the internet, it was paper-work. Thousands of pieces of paper. Getting all the clipboards together, transcribing notes and recording them in a ledger. Doing calculations by hand to work out how many people did what, when, and how that impacted the company’s net income. Today, all of that calculation

happens in the blink of an eye. Anyone can do it, using web-based software products like Google Analytics, Adobe Analytics, Mixpanel, Hotjar, or Amplitude.

These tools can give a wide manner of different insights: which ads or channels are driving traffic online, which pages are most effective at persuading users to take actions, the step in a series of pages at which users give up because they're confused or frustrated, and more. All of these insights are then looped back into the design process, where changes are made to the product to boost *conversion rates*: the proportion of people who complete an action compared to those who do not.

EASIER A/B TESTING

A/B testing⁴ was first used commercially in the early 20th century, but in those days it was an awkward, painstaking process.⁵ You could do it with newspaper ads: you'd run one version of your ad with a coupon, and another version with a different coupon. The version that won was the one that got the most coupons used. In those days, all the work was done by people; coupons delivered back to the agency were manually sorted and tallied by admin staff. It was a load of work and, of course, if your business wasn't all about advertising general consumer products to the masses, you were stuck.

The limitations of the physical world mean you can't apply the same kind of A/B testing to physical products and services as to digital without a great deal of cost and uncertainty. For example, if you have a shop on the high street, you can't change the store layout from one customer to the next. Perhaps if you were Cobb from the film *Inception*, you'd be able to click your fingers and rearrange your shop floor at a whim. In the digital world, *Inception*-like remodelling is trivially easy. You can make two versions of a page or feature and easily find out which performs better. For example, version A of a page might say '20 other people are looking at this item', while version B of the same page might say 'Only 2 items left in stock'. Your A/B testing software then deploys version A to a random sample of users and version B to another. After the test is complete, your A/B testing software will

automatically calculate statistics for you, telling you if either of the designs performed significantly better than the other on the measured conversion rate (purchases completed, for instance). You don't even need to understand the statistics, as the results are usually dumbed down into simple sentences for you. No magic, cement, bricks or PhD needed. In fact, creating an A/B test today is as simple as signing up to a product like VWO or Optimizely free of charge and filling in a few forms.

A/B testing doesn't judge whether a particular design is actually better or worse for the user – it just provides statistics as to whether design A or B performed better on your chosen metric. This means A/B testing opens a door towards deceptive patterns, because when a business tests a deceptive pattern against a more neutral pattern, typically it's found to perform better on the chosen metric. Why? Because tricking or trapping users can be more effective than persuading them; and also because persuasion is frequently *combined with* deception, which means the overall page has two shots at capturing the user. It can start out by trying to persuade the user to complete the desired action. Then, if the user isn't successfully persuaded, the deceptive pattern has a chance to get them to complete the desired action through nefarious means. Imagine some persuasive content followed by a preselected checkbox, for example. Some users will be persuaded by the content and will be happy with the default. Others won't be persuaded and also won't notice the preselected checkbox, so they'll end up being tricked into opting into something they didn't want.

When a deceptive pattern wins an A/B test, it's often a direct source of revenue, with statistics to prove its effectiveness. In a metrics-driven environment, it can be very hard for employees to push back against this and encourage a more user-friendly – but less profitable – approach.

COPYCAT DESIGN

It was Oscar Wilde who said, 'Imitation is the sincerest form of flattery that mediocrity can pay to greatness'. Some tech companies have been

very successful in driving up conversion rates by using deceptive patterns. In response, others have copied them. This isn't at all surprising. If you saw a competitor successfully making money for years without any legal or regulatory consequences, then why wouldn't you copy them?

CHAPTER 4

FROM HOMO ECONOMICUS TO HOMO MANIPULABLE

To understand deceptive patterns, we need to understand some concepts from the field of economics. For a long time, economists believed humans were perfect information-processing machines – able to consume, understand and reason with all the information provided to them at all times. They called this idea ‘homo economicus’. If you think about the number of mistakes we all make in our daily lives, you’ll know this is a really daft idea. Still, it’s understandable. And economists needed to start somewhere, and they also needed to start with a relatively simple model of how humans behave, otherwise the maths gets really complicated.

It’s only relatively recently – in the late 20th century – that economists have updated their views. It was considered groundbreaking when Herbert Simon introduced the idea of ‘bounded rationality’.¹ He posited that ‘both the knowledge and the computational power of the decision maker are severely limited’ and ‘we must distinguish between the real world and the actor’s perception of it and reasoning about it’. In other words, we can only remember a certain amount of stuff before we start forgetting; we can only do a certain level of mental arithmetic before we get it wrong; and we can only read so much complex text before we become fatigued and start to misunderstand things.

To be even more reductionist, bounded rationality means we muddle through life doing our best with limited faculties. As someone who once fell down the stairs at night because I had forgotten that I'd moved house, I can attest to that.

More recently, behavioural economics has greatly extended the idea of bounded rationality. Richard Thaler is considered one of the founders of behavioural economics, and he won the Nobel prize in 2017 for 'incorporating psychologically realistic assumptions into analyses of economic decision-making'.² It turns out that understanding the ways in which people can do dumb things is really useful for economic modelling. Particularly when it comes to understanding the causes of the common mistakes we all make.³

'Real people have trouble with long division if they don't have a calculator, sometimes forget their spouse's birthday [...]. They are not homo economicus; they are homo sapiens.'

—Thaler and Sunstein (2008)

Physically, our bodies have lots of common flaws. For example, the trachea and oesophagus are very close to each other. Most of us are familiar with the dangers of accidental choking. Knowing that flaw and sharing the knowledge has helped humanity a great deal. The same applies to human reasoning and decision-making. If we can understand ourselves better, the more likely it is that we'll be in a position to overcome our weaknesses.

Most psychology researchers and theorists are motivated by this honourable goal: improving the human condition. There's even a branch of applied psychology – human factors and ergonomics – which aims to 'reduce human error, increase productivity, and enhance safety and comfort'.⁴ In a nutshell, the aim is to understand how the human mind works, and then use those insights to help people make better decisions.

Unfortunately, not everyone is motivated by kindness. Some see human weakness as a commercial opportunity. Instead of thinking of humans as *homo economicus*, it is perhaps more useful to think of us as ‘*homo manipulable*’: imperfect and vulnerable to control by others in ways we may not even notice.⁵

To recap, this chapter has explored the rise of deceptive patterns in the digital world and the reasons behind their ubiquity. Several key factors are identified as contributing to the proliferation of deceptive patterns, including the emergence of a metrics-driven culture, the ease of tracking and data processing, the widespread use of A/B testing, and the prevalence of copycat design in the tech industry. Over the past few decades, well-intentioned academic research has revealed weaknesses in human reasoning and decision-making. Today, these insights are used to manipulate users for profit, which is a far cry from the original intent of the research.

PART TWO

EXPLOITATIVE STRATEGIES

There are lots of different ways you can consider the underlying psychology and principles behind deceptive patterns. A good starting point is to think of them as the result of an exploitative business strategy. In other words, instead of a business thinking of its users as partners who should be cooperated with to reach mutual success ('Their success is our success'), the business thinks of its users as a commodity to be exploited ('Their weakness is our opportunity'). Another aspect of the exploitative mindset is the business's attitude towards law: rather than seeing it as a system to be respected, the law might be seen as a game to be played, where loopholes and grey areas can be identified and exploited for profit.

Exploitative		Cooperative
"Your weakness is our opportunity"		"Your success is our success"
Attitude to users	User as commodity	User as human
	Vunlerabilities exploited for profit	Vulnerablites supported with care
Attitude to the law	Law as game to be played	Law as system to be respected
	Loopholes as growth opportunity	Loopholes as pitfalls to avoid
Result	Deceptive patterns	User-centred patterns

A comparison between exploitative and cooperative design strategies.

If we look at it in a simplistic way, exploitative strategies are often going to be more effective than cooperative strategies because they sidestep the need to let users make an informed choice. It's a bit like wondering whether a fishing net is going to be more effective than just *asking* fish to jump into your boat. The fishing net is a trap, similar to a deceptive pattern. If you impede a user's ability to make an informed choice, or if you hinder their decision-making by hiding facts or by giving misleading information, then you effectively capture or lock in the user against their will (though they may not realise it at the time owing to a lack of clearly stated information).

Generally, businesses do not admit to themselves that they are using exploitative strategies. If a business focuses on growth and measured outcomes, it can slip into an exploitative mindset without realising it. Euphemisms are also very common in businesses (for example, a subscription that automatically renews without an email reminder might be glossed as 'We are helping users enjoy an uninterrupted service'), and the true consequences of a design decision may be far away from the people who implement it. In large tech companies, customer service teams are often outsourced overseas, far from headquarters where the decision-making happens. When users are presented as charts and data dashboards in executive meetings, their humanity is stripped away, and it's easy to slip into thinking of them as just numbers, a commodity to be processed and from which value is extracted.

The best way to explain deceptive patterns is to start by looking at the exploitative design strategies – so you come to understand the theory, principles and goals – and then look at the result of the strategies, so you can then have a good basis for understanding the specific types and examples of deceptive patterns in the wild.

Professor Colin M. Gray and his team at Purdue University's UXP2 Lab were among the first researchers to look closely at the exploitative design strategies that lead to deceptive patterns.¹ Expanding on their work, I present eight types of exploitative design strategy in this chapter. These are summarised below.

- **Exploiting perceptual vulnerabilities:** Before a human can reason about information, they have to perceive it first. Since human perception is not perfect, the shortcomings can be exploited to hide information, e.g. low contrast, small text.
- **Exploiting vulnerabilities in comprehension:** Humans have limits to literacy, numeracy, critical thinking and memory. An exploitative designer can make something more complicated than it needs to be, e.g. the use of verbose terms and conditions.
- **Exploiting vulnerabilities in decision-making:** Cognitive biases are systematic errors in reasoning that all humans tend to make. They can be exploited to interfere with decision-making, e.g. a preselected checkbox can take advantage of the default effect.
- **Exploiting expectations:** Helpful design involves employing standards to make a product predictable for users. These standards can be subverted to trick users e.g. making an 'X' button mean 'yes' instead of 'no'.
- **Resource depletion and pressure:** Humans have a limited supply of attention, energy and time. Once these resources become depleted, users may give up; they may feel pressure; and they may become fatigued and vulnerable to other tricks. e.g. cookie consent dialogs often require extreme effort to opt out, wearing users down until they give in.
- **Forcing and blocking:** 'Forcing' involves putting a mandatory step in front of the action the user wants to complete, which they cannot decline, e.g. mandatory registration in order to complete a purchase. 'Blocking' involves the outright removal of a feature, e.g. preventing the user from exporting their own data.
- **Exploiting emotional vulnerabilities:** Humans do not like to experience uncomfortable emotions like guilt, shame, fear or regret, and will often take measures to avoid them. e.g. to decline an offer for a fitness course, the user must click 'No thanks, I want to be unhealthy.'

- **Exploiting addiction:** Humans are prone to addiction, where a habit develops harmful outcomes and becomes difficult to give up. This involves a cycle of behaviour that can be intensified through design techniques like infinite scroll or autoplay.

CHAPTER 5

EXPLOITING PERCEPTUAL VULNERABILITIES

Before a human can reason about information, they first have to perceive it. With so much of our lives being online and on-screen, it is useful to consider how visual perception works.

Although it's tempting to consider healthy human eyes as perfect high-definition cameras, they are actually very different.¹ In fact, much of what we visually perceive is fabricated by the perceptual systems in the brain, and our eyes provide highly incomplete information. For example, the human eye has a physical blind spot at the back, where the optic nerve connects the eyeball to the brain. People with normal vision do not see any blind spot, despite it being present at all times. It's filled in by our visual cortex.² In simple terms, the human visual perception system guesses at what should be present and fabricates it. Or in other words, our brains are making it up as we go along.

In the same way, the middle of the retina contains sensors called cones that enable us to see in colour. Around the periphery of the retina we mainly have a different type of receptor called rods, which provide non-colour vision, and work better in low light. Yet we do not perceive any variation in the colour of what we see from the centre to the periphery of our field of vision. The human visual cortex does an enor-

mous amount of ‘guesswork’ to fill in the gaps, making an inconsistent data source seem utterly full colour and high definition.

What’s more, when a person with normal vision looks at something, they usually perceive a steady, fixed scene. However, the human eyeball typically moves around a great deal. When we read something or scan our surroundings, our eyes rapidly flick from here to there and back again, taking in all kinds of pieces of information in addition to what we’re focused on. The fast movements are called *saccades*, and they last somewhere between 20 and 200 milliseconds. They’re interspersed by *fixations*, when the eye stops, briefly, for 50 to 600 milliseconds. Yet we don’t get motion sickness from it – we don’t notice it at all.

To summarise, what you ‘see’ as a human is not reality, but an internal representation of reality involving imperfect sensors (our eyes) and an enormous amount of internal processing that uses amazing guesswork to fill in the gaps. This means that the entire visual system is exploitable, making it easy to hide things. In other words, camouflage.

A famous example of camouflage in nature is the lime hawk-moth (*Mimas tiliae*),³ which has evolved to blend into its environment using colours and visual contrast to mimic its background and to create false edges that disrupt its shape, thereby avoiding visual detection by predators when positioned on a lime tree or similar vegetation.⁴



Visual camouflage used by a lime hawk-moth.

With apps and websites, exploitative designers frequently use similar techniques, by manipulating text colour contrast and size.

The interesting thing about colour contrast is that it is straightforward to calculate.⁵ You capture the hex codes for the text foreground and background colours, then enter the values into a calculator tool.⁶ There is an internationally recognised standard for minimum colour contrast: the W3C's Web Content Accessibility Guidelines (WCAG 2.1). It has three levels. The middle level, 'AA', is widely recognised as the baseline to aim for.⁷ This means you can use a colour contrast calculator tool to instantly work out if a piece of text on a page meets the baseline recommended level for colour contrast.

One trick to watch out for is differences in text contrast on a page. If most of the text on a page is relatively high-contrast and one bit of text is relatively low-contrast, this could make readers less likely to notice or pay attention to the lower-contrast text, even if it is AA-level compliant. Readers often interpret colour contrast as a signal of what

they should read versus what they can safely ignore (i.e. ‘this pale grey text can’t be very important - if it was, they’d have made it more prominent’).

One of the first cases I worked on as an expert witness was *Arena vs Intuit Inc.*⁸ In 2019, a law firm called Stueve Siegel Hanson approached me and asked me to review the account creation and sign-in process relating to Intuit’s TurboTax products. A screenshot of the sign-up page is shown below. See if you can spot any issues, based on what we’ve talked about in this section.

The screenshot shows the TurboTax sign-up page. At the top, the Intuit logo is centered, with logos for TurboTax, QuickBooks, and Mint below it. The main heading reads "One Account. Everything Intuit." followed by a subtext "One account for everything Intuit, including TurboTax. [Learn more](#)". The form consists of four input fields: "Email address", "User ID", "Phone (recommended)", and "Password". The "Password" field has a toggle icon on the right. Below the fields is a blue "Create Account" button with a lock icon. At the bottom, there is a disclaimer: "By clicking Create Account, you agree to the [TurboTax Terms of Service](#), [Turbo Terms of Use](#) and have read and acknowledge our [Privacy Statement](#)." and a footer for Invisible reCAPTCHA by Google with links to [Privacy Policy](#) and [Terms of Use](#).

intuit

turbotax qb quickbooks mint

One Account.
Everything Intuit.

One account for everything Intuit, including
TurboTax. [Learn more](#)

Email address

User ID

Phone (recommended)

Standard call, messaging or data rates may apply.

Password

Create Account

By clicking Create Account, you agree to the
[TurboTax Terms of Service](#), [Turbo Terms of Use](#)
and have read and acknowledge our [Privacy Statement](#).

Invisible reCAPTCHA by Google [Privacy Policy](#) and [Terms of Use](#).

Screenshot of the TurboTax sign-up page in November 2019.

You might not realise it by looking at the screenshot in the figure above, but if you clicked the 'Create Account' button, you would be agreeing to binding arbitration. In other words, you would be unable to take Intuit to court. To find the information about arbitration, you are expected to notice and read the text below the big blue button ('By clicking Create Account...').

In my analysis I found that the colour contrast of that text was lower than most of the other text on the page, and the font size was smaller too. I can't write too much since a good deal of my report is confidential, but the key point here is that the judge agreed with this analysis. To quote:

'...both the notice and the hyperlinks therein are in the lightest font on the entire sign-in screen [...] The Court finds that a reasonable consumer would be less likely to notice text in a significantly fainter font than other text on the same page.'

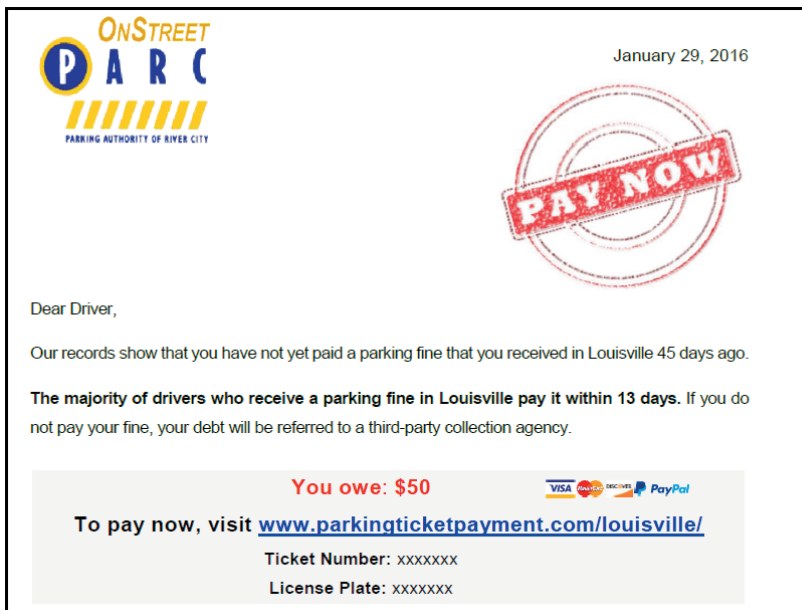
—CHARLES R BREYER, UNITED STATES
DISTRICT JUDGE, 12 MARCH 2020

During the case, Intuit were required to disclose their analytics data, which showed that less than 0.55% of users actually clicked on the relevant hyperlink during a four-month period in 2019⁹. This is compelling evidence, and also a reminder that internal company data or documents can end up being revealed in a lawsuit, and it's sometimes the job of an expert witness to suggest what to ask for.

To summarise, small low-contrast text is an effective way to hide content on a page, and to prevent users from subsequently comprehending it and making an informed decision. Like many other exploitative strategies, manipulation of perception may be illegal in some jurisdictions, as Judge Breyer found above, depending on the way you use it.¹⁰

Another more brazen approach to exploiting perception is to remove something entirely from the user's perceptual field. If you don't want people to comprehend something, you just don't show it on the page, and you put it behind links or buttons that allude to something else. This is a very common practice in cookie consent dialogs – where the first thing the user sees gives no clue that there may be a button somewhere that allows them to reject all forms of tracking. In 2020, Nouewens et al. carried out a research study to measure the impact of this.¹¹ Forty participants took part in an online field experiment. Results showed that removing the 'Reject all' button from the first page of a consent pop-up increased consent by up to 23%.

In another study, the UK government's Behavioural Insights Team (BIT, aka 'the nudge unit') worked with an Australian government department to improve payment rates for fines, debts and taxes¹². They sent two different letters to 48,445 people. In one letter, they featured a large red 'pay now' stamp on the letter, shown below. The other letter didn't have this stamp.¹³



Example of a letter with a 'pay now' stamp from a similar study by BIT.

They found that the letter with the 'pay now' stamp delivered a 3.1 percentage point increase (14.7% payment rate without the 'pay now' stamp; 17.8% payment rate with it).¹⁴ We can look at that figure the other way around: the letter *without* the 'pay now' stamp delivered far *fewer* payments. So the simple act of removing a call to action is very effective. If someone doesn't perceive something – such as the need to act, or a reason to do something – they are less likely to think about it. If they don't think about it, it won't influence their decision-making.

There are other ways to manipulate perception that I should mention before we move on to the next type of strategy. Most commonly, exploitative designers employ clutter and noise – subverting common graphic design principles like white space, repetition, alignment and proximity. (If you're not familiar with these principles, any introductory textbook can provide you with a beginners guide.¹⁵) This serves to create a sort of smokescreen, making elements of the page harder to see, read and scan: playing into the exploitation of expectations strategy and the exploitation of vulnerabilities in comprehension strategy.

CHAPTER 6

EXPLOITING

VULNERABILITIES IN

COMPREHENSION

LITERACY, NUMERACY AND PROBLEM SOLVING

In 2013, a huge worldwide study called the Program for the International Assessment of Adult Competencies (PIAAC) was published, involving a quarter of a million participants in 33 different countries.¹ It looked at literacy, numeracy and problem-solving proficiencies across the world. The following summary is just for the United States, though the picture is fairly similar in many countries. According to the 2013 PIAAC findings:²

- 30% of adults in the US are likely to have difficulty sorting through emails and organising them in folders provided for them.
- 20% of adults in the US are unlikely to find the name of a congressperson with a summary information sheet that lists the district, name, year and place of birth.
- 30% of adults in the US are unlikely to be able to calculate the total cost of daily car rental when provided with miles driven that day, cost per day and cost per mile.

- 16% of adults in the US are digitally illiterate, and cannot use a computer to find a recipe, make a retail purchase or file taxes online.

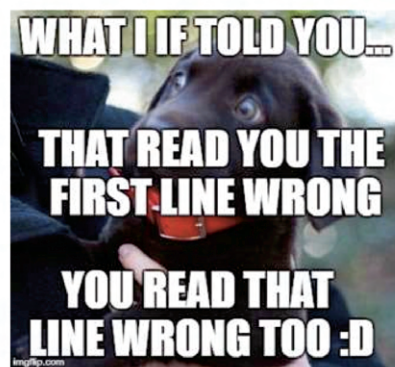
As you can see, low literacy and numeracy is very common. With an exploitative mindset, this presents an exploitable vulnerability. If a business wants to hide unfair or unappealing aspects of a transaction, it can do so through the use of complex language or complex numerical content. With this in mind, it's interesting to consider the writing style used on public service websites – plain language, short sentences, and enormous efforts taken for comprehension for all citizens – versus the writing style used in more exploitative products like crypto trading apps, where impenetrable technical terms are used extensively, very little is explained, and the user is enabled to make all kinds of risky trades with minimal safeguards or education.

HOW SCAN READING CAN BE USED TO MANIPULATE PEOPLE

When we read, we don't usually read every word on every page. Not unless we're studying really hard or working our way through something we're enjoying, like a novel for example. Take a look at this:³



And you will read this at the end

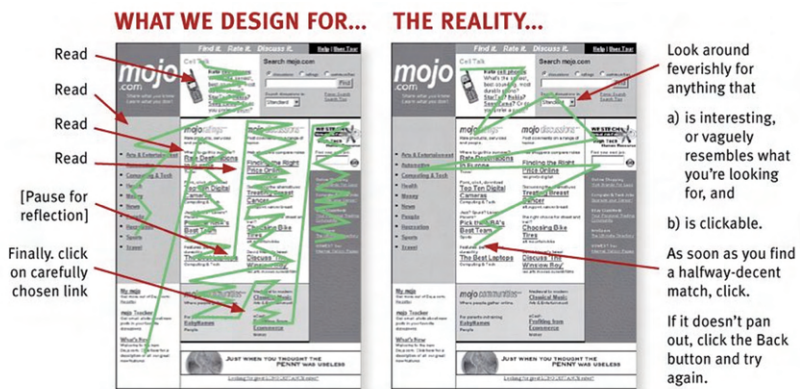


A demonstration of how human visual perception can be manipulated.

On the left, you can see we tend to let visual hierarchy determine the order in which we read things. We've learned it's a good idea to read the big, prominent things first and smaller things later. On the right, you can see how we glance at content and make educated guesses about individual words to save time. This isn't something we're born with. It's a technique called 'scan reading', which we pick up naturally as we get better at reading. Similarly, good writers and page designers learn how to design for scan reading, to help people to do it more efficiently.

Steve Krug published *Don't Make Me Think* in 2000. It's now in its third edition, with over 350,000 copies in print. This book is highly regarded in the UX design industry, as it puts forward a clear explanation for the concept of scan reading by people who are using screens.

Let me show you two more images. The one on the left explains what we naively might assume is a natural way to read information. In theory, we'd expect readers to take in each successive word, thoroughly paying attention to every element of the design. This naive view of human information-seeking behaviour is similar to the concept of homo economicus from traditional economics – the idea that humans have a limitless supply of attention, energy and critical thinking skills, so we can brute-force our way through any body of content by reading every word on a page in sequence.



A depiction of scan reading behaviour (Krug, 2006)

However, Krug argues that while authors might hope that people read every word on every page in a highly attentive and rational manner, the reality is rather different. In real life, most of us take a ‘billboard going by at 60 miles an hour’ approach⁴ when there’s this much information presented to us.⁵

Krug argues that users tend to ‘glance at each new page, scan some of the text, and click on the first link that catches their interest or vaguely resembles the thing they’re looking for. There are usually large parts of the page that they don’t even look at.’ He goes on to explain that we’ve been trained to scan-read from an early age, flicking through newspapers and magazines, for example, or reference books, as we try to narrow down many choices and find just the parts we’re interested in.

Other researchers found more evidence. Back in 1997, Morkes and Nielsen did a quantitative empirical study in which 51 participants tested five variations of a website, each one with a different style of writing:⁶

1. A promotional writing style – full of ‘marketese’
2. A scannable writing style – intended to encourage scan reading
3. A concise writing style – succinct content
4. An objective writing style – not using promotional language
5. A combined concise, scannable and objective writing style

Each person was given a series of tasks, generally involving looking for the answer to a simple question. The amount of time they took was recorded, as were any errors they made. The findings showed that people performed worse on the promotional style pages, but significantly better with the scannable and concise styles.

This research demonstrated what we might have divined naturally: a writing style has an impact on users’ ability to read and understand information. If users read every word on every page in a systematic way, these differences wouldn’t have been seen. In a subsequent article, Nielsen (1997) addressed the question, ‘How do users read on the web?’ with a two-word answer: ‘They don’t.’ He went on: ‘People

rarely read web pages word by word; instead, they scan the page, picking out individual words and sentences.⁷

Understanding how people read is vital if you want to design web pages or app screens that work effectively, or – conversely – if you want to create deceptive patterns.

Eye-tracking research is another useful source of insights about reading behaviour. In 2014, Pernice, Whitenon and Nielsen ran an eye tracking study with over 300 participants.⁸ In one exercise, people were asked to use a search engine and find some specific information. Eye-tracking technology followed their progress, monitoring how they fixated on the page: 17% of the time, people looked at only one result before clicking onto the next page. They didn't fix their gaze anywhere else. Or, in other words, they picked the first result that seemed 'good enough' to save effort, rather than systematically reading every result on the page. This is a demonstration of an *information foraging strategy*, a technique that was first defined in 1999 by Pirolli and Card, who noticed similarities between animal food foraging strategies and the way in which humans search for information online.⁹ When an animal forages for food, it cannot search everywhere or it may die from starvation, so it must use a 'good enough' strategy that provides the most benefit for the lowest cost. Broadly speaking, information foraging can be considered a kind of goal-directed scan-reading strategy.

Generally, scan reading and information foraging can be a pretty effective way of saving ourselves time and energy. But it is only effective in a predictable, trusted environment in which the designer has your best interests in mind. If a designer wants to deceive you, they can take advantage of scan-reading behaviour by hiding pertinent information where you don't expect it, or by using misleading headings and visual hierarchy, among other things.

MISLEADING INFORMATION AND FALSE BELIEFS

If a business publishes misleading information, this can lead users to make decisions that are not in their best interests. This can range from outright lies (fraudulent claims) to ambiguous or manipulative

language and design that encourages the user towards a false belief. For example, a business might exploit scan reading by making the headings, links and buttons on a page appear to say one thing, while the body text, if read word-for-word, says another. Similarly, offers can be priced in a manner that requires considerable mental arithmetic and short-term memory to compare properly. If the user is not capable of this task, they might end up with a bad deal that harms them financially. The FTC lists ‘false beliefs’ as one of the top harms posed to consumers by deceptive patterns¹⁰. In a 2021 study involving 3,777 participants, researchers Luguri and Strahilevitz found that ‘hidden information’ doubled the acceptance rates for a product offer, as compared to a neutral design. In other words, participants formed false beliefs about an offer because facts were hidden away from view, and this had a substantial effect on their decisions.¹¹

CHAPTER 7

EXPLOITING VULNERABILITIES IN DECISION-MAKING

If you think of the stream of information that enters your mind, you first have to perceive it, and then you have to comprehend it. I've explained how weaknesses in both of these areas can be exploited. After perception and comprehension occur, we then need to engage in critical thinking, or what cognitive psychologists tend to call 'judgement and decision-making' which can also be exploited for commercial gain.¹ To quote whistleblower Christopher Wylie from his book *Mindf*ck*:²

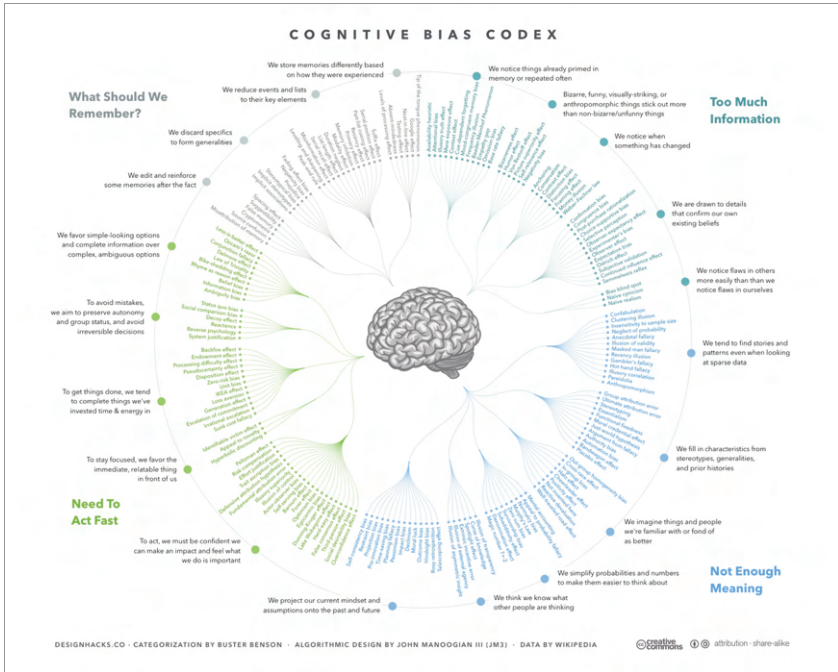
'The goal in hacking is to find a weak point in a system and then exploit that vulnerability. In psychological warfare, the weak points are flaws in how people think. If you're trying to hack a person's mind, you need to identify cognitive biases and then exploit them.'

— Christopher Wylie, p.63

A cognitive bias is a mental shortcut that tends to cause a systematic error in judgement and decision-making. Humans fall foul of these biases rather predictably, which led economist Dan Ariely to describe

human behaviour as ‘predictably irrational’.³ Despite their shortcomings, cognitive biases are also believed to provide benefits because they provide shortcuts, ways to avoid effortful work in order to save time and energy for other more important matters. Cognitive scientist Aaron Sloman describes this as ‘productive laziness’ and explains, ‘a chess champion who wins by working through all the possible sequences of moves several steps ahead and choosing the optimal one is not as intelligent as the player who avoids explicitly examining so many cases’.⁴ Sloman wrote this in 1988 – no doubt he would happily refer to the web instead of chess if he were to write it today. No sensible human would read every result on Google, or every product listing on Amazon before choosing which item to click. Shortcuts are necessary to cope, so today we rely on cognitive biases more than ever, because we simply cannot process all the information we receive in detail.

There are thousands of research papers and well over one hundred types of cognitive biases proposed, though not all are considered rigorously researched. You can get a sense of the range and types by looking at the cognitive bias codex (though it’s best viewed on a large screen given the information density).⁵



The cognitive bias codex (Manoogian & Benson, 2016)

Research on cognitive biases started to become well known in the early 2000s, entering the realms of pop psychology, business and design textbooks. The tech industry latched onto this with a great deal of enthusiasm. Some authors were very direct about the purpose of their work. In the introduction of his book *Influence*, Robert Cialdini refers to his area of work as ‘the psychology of compliance’ (that is, submission to demands of others) and he describes his key principles as ‘six universal weapons of influence’.⁶ In the book *Hooked*, the author Nir Eyal promotes a ‘habit-forming’ behavioural model that is nearly identical to Natasha Dow Schüll’s model of ‘ludic loops’ – except Dow Schüll describes her model as ‘addiction by design’ and presents harrowing accounts of lives destroyed by gambling.⁷ Eyal is careful to avoid the word ‘addiction’, but the connection is obvious.

Today, numerous websites and blogs provide guides on how to exploit cognitive biases for profit; for example, the company Convertize provides a library of cognitive biases that it cheerfully recommends as

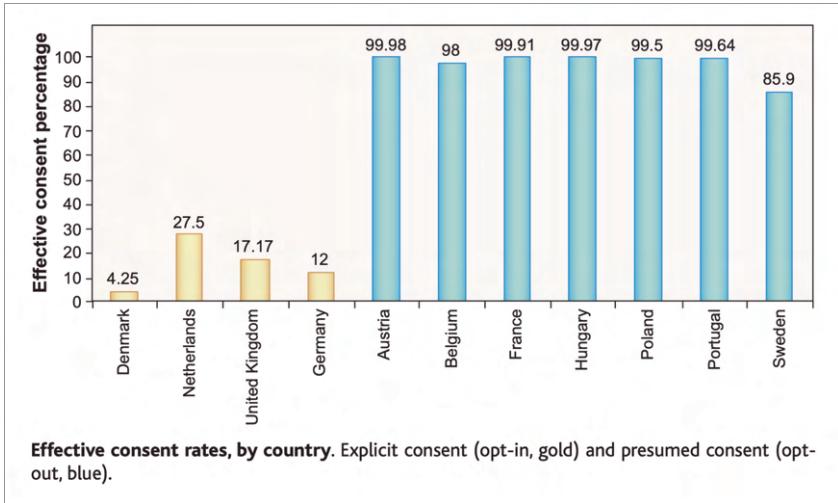
‘A/B Testing Ideas Based On Neuromarketing’, without any mention of negative consequences for the end user, such as being tricked or trapped into unwanted transactions or contracts.⁸

There’s also lots of content available about cognitive biases and persuasion that proposes use in a non-exploitative manner – but it’s a very short hop from ‘use this bias to persuade in a transparent and helpful way’ to ‘use this bias to see what happens in your next A/B test’. After all, as soon as a design is tested and has statistical evidence proving it to be more profitable than the other designs, it’s very likely to be adopted by the business with little further discussion, regardless of whether users truly understand the consequences of their actions.

DEFAULT EFFECT

The default effect is a psychological phenomenon where people tend to stick with the status quo and choose the option presented to them as the default. It’s a bias that’s been studied in many different contexts, from consumer decisions to public policy. Businesses know that people are more likely to stick with the default option, so they often define the default to be favourable to the business in some way, typically through a preselected checkbox or radio button.

One of the most famous studies on the default effect was carried out by researchers Eric J Johnson and Daniel Goldstein in the 2003 paper ‘Do Defaults Save Lives?’ They looked at organ donation consent rates in different countries, and they compared the countries in which users are opted out by default (shown on the left) versus countries in which users are opted in by default (shown on the right).



Effective consent rates by country, from Johnson & Goldstein (2003).

As you can see, the difference in consent rates was enormous. A number of things are believed to drive the power of the default effect:

- **Awareness:** for a user to change the default, they first have to become aware that it is possible to do so. (This harks back to the earlier section on exploitation of perceptual vulnerabilities.)
- **Effort:** for the user to change from the default, they have to do something; in this case it involves finding and completing the correct government form. It is possible that citizens might intend to change their choice from the default, but not have time or energy to do so.
- **Authority bias and social proof:** the default effect can be combined with other cognitive biases. For example, the default may be presented as the correct thing to do by a figure of authority (a doctor, for example). Alternatively, it may be portrayed as the thing that everyone else is doing (social proof). These are both known to be powerful cognitive biases in their own right.

In the book *Misbehaving* Richard Thaler did some follow up research, looking at *true* organ donation rates as opposed to *presumed* consent rates.⁹ He found that while that presuming consent may appear to work on paper, when people die in hospitals the staff will typically ask the family whether the organs should be donated. At that point the presumed consent frequently gets discarded as there is no record of the individual's actual choice. Thaler concluded that 'mandated consent' was a better policy, forcing citizens to make an explicit choice when they renew their driving licence.

The default effect has also been studied in the context of privacy and cookie consent dialogs. A large-scale study conducted by SERNAC, the Chilean consumer protection agency, provides compelling evidence.¹⁰ Over 70,000 participants were presented with different cookie consent interfaces. In one of the interfaces, participants were presented with cookie tracking as opted-in by default, while another presented it as opted-out by default. The opted-out version increased the rate of users rejecting cookies by 86 percentage points.

As you can see from the evidence, the default effect is easy to employ and is very powerful. It is often used by businesses in an exploitative way: to presume user consent for decisions where users might prefer to opt out, if they only knew the true nature of the decision they were being presented with, and were given an explicit choice.

ANCHORING AND FRAMING

The anchoring effect cognitive bias is a psychological phenomenon where individuals rely too heavily on the first piece of information they receive (the anchor) when making decisions. For example, Tversky and Kahneman (1974) conducted a study in which participants were asked to estimate the percentage of African countries in the United Nations.¹¹ They were first given a random percentage (an anchor), then asked if their estimate would be higher or lower, and then finally asked to provide their estimated figure. The results showed that the estimates of participants were significantly influenced by the anchor they were given: those given a higher anchor estimated a

higher number, and those given a lower anchor estimated a lower number. This insight is frequently used by marketers in an exploitative manner when pricing consumer products – for example, where an initial price is created to be artificially high so that a discount can be presented, giving a sense of value for money.

Framing is a similar cognitive bias where individuals rely too heavily on the way information is presented rather than on the underlying facts. In 1981, Tversky and Kahneman carried out an experiment in which they gave participants a scenario relating to a hypothetical disease, and were given two treatment programmes to choose from.¹² Depending on their experimental group, the outcomes of the treatment programmes were framed either positively: ‘X people will be saved’; or negatively: ‘Y people will die’. They found that the framing had a pronounced effect on participants’ choices, even though the underlying facts were identical in both cases.

In the book *Predictably Irrational*, Dan Ariely reported a study that demonstrates the manipulative power of this type of cognitive bias.¹³ He created two different fictional designs of *The Economist* magazine’s subscription page, and presented them to 200 students (100 per design), asking them to pick their preferred subscription type. Unknown to the participants, one of the designs contained a trick (design A, below), intended to get participants to perceive the combined print and web subscription as better value. It involved providing an extra ‘decoy’ subscription: the print magazine on its own for the same price as the print and web subscription. As you can see in the figure below, the presence of the decoy print subscription in design A caused the print and web subscription to be selected much more frequently (84% selected) than when it was omitted in design B (32% selected).

Design A: most participants selected the print & web subscription because the identically priced print-only subscription served as a decoy.

<input type="checkbox"/> Economist.com subscription - US \$59.00 One-year subscription to Economist.com. Includes online access to all articles from <i>The Economist</i> since 1997.	Selected by 16/100
<input type="checkbox"/> Print subscription - US \$125.00 One-year subscription to the print edition of <i>The Economist</i> .	Selected by 0/100
<input type="checkbox"/> Print & web subscription - US \$125.00 One-year subscription to the print edition of <i>The Economist</i> and online access to all articles from <i>The Economist</i> since 1997.	Selected by 84/100

Design B: when the print subscription was removed, fewer participants selected the print & web subscription.

<input type="checkbox"/> Economist.com subscription - US \$59.00 One-year subscription to Economist.com. Includes online access to all articles from <i>The Economist</i> since 1997.	Selected by 68/100
<input type="checkbox"/> Print & web subscription - US \$125.00 One-year subscription to the print edition of <i>The Economist</i> and online access to all articles from <i>The Economist</i> since 1997.	Selected by 32/100

Dan Airley's Economist magazine study, where the presence of a decoy option influenced participants' decision-making.

SOCIAL PROOF

The social proof cognitive bias is a phenomenon in which individuals are likely to conform to the behaviour of others. It's also known as the 'bandwagon effect', 'groupthink' or the 'herd effect'. To put it another way, if we see that numerous other people perceive something as valuable, we are likely to believe that they are correct. This is a shortcut that allows us to avoid the hard work of carrying out a critical evaluation of our own.

In 2014, a group of researchers working with HMRC tested the impact of social proof in a large-scale experiment.¹⁴ They designed five

different tax bill reminder letters, each with a different message, shown in the table below. They sent these letters to a random selection of 100,000 UK taxpayers, and tracked the response rate (which they measured as a successful payment of the tax bill within 23 days).

No.	Message	Response rate
1	"Nine out of ten people pay their tax on time"	1.30%
2	"Nine out of ten people in the UK pay their tax on time"	2.10%
3	"Nine out of ten people in the UK pay their tax on time. You are currently in the very small minority of people who have not paid us yet"	5.10%
4	"Paying tax means we all gain from vital public services like the NHS, roads, and schools"	1.60%
5	"Not paying tax means we all lose out on vital public services like the NHS, roads, and schools"	1.60%


Findings from HMRC tax letter study (Hallsworth et al., 2017).

As you can see, messages 1, 2 and 3 used different styles of social proof, while messages 4 and 5 did not. Message 3 employed the most aggressive social proof phrasing and it was by far the best performing. This was a big win for HMRC, and timely tax payments benefit the country as a whole. Of course, there's nothing exploitative about this example – accurate and true social proof information is constructive and helpful. However, it can become exploitative when the information is tampered with in some way, and the user is purposefully not informed about what's going on.

Online, social proof is typically presented as reviews, case studies, testimonials and data (ratings or 'likes'). For example, consider a testimonial. If it is completely fabricated by the company, then that's just false advertising – fraud, plain and simple. Similarly, if it's provided by a real user but they were paid to write something positive, then that's fraudulent too.


But what if it's real, and the user was paid to give an honest and unbiased review? Incentivisation creates a grey area in which exploitative practices can be hidden. For example, what kind of payment was the

reviewer given? Was the payment proportional to the service provided? Did the company imply that future employment as a reviewer might be conditional on a positive review this time? Did the reviewer give a positive review because of the incentive, even though they were not asked to? We all know from personal experience that if we receive a gift or a big discount we will be less critical of its shortcomings than if we had paid for it ourselves at full price. So, incentivised reviews should always be labelled with a disclosure – the user needs to be told that the review was paid for. However, the problem with disclosures is that they can be ambiguous. Take this Amazon UK review for an airfryer:¹⁵

 Steven E. **VINE VOICE**

★★★★★

Brilliant

Reviewed in the United Kingdom  on 17 July 2021

Verified Purchase

The ninja dual zone Air fryer has a huge 7.6 L total capacity. There is two individual baskets which are 3.8 L each. It has six programs, which are Air fry, roast, bake, dehydrate, Max crisp and reheat. If you are only going to use one of the baskets then zone one is the default zone to cook with, but what I particularly like about this is that for instance you can air fry at 200° in zone one for 30 minutes and you can roast in zone two at 180° for 20 minutes at the same time by pressing sync and then start. This will start zone one and when it reaches the 20 minute timer zone two will come off of hold and start roasting. If you remove the basket such as for chips to give them a shake, then the timer will pause until The basket is re-inserted and if both zones are cooking it will pause both timers. The other option that I really like which is nice and easy and can maximise your output if you're cooking the same thing, For example if you had to Air fryer 1.5 kg of chips. Then place your chips in both baskets and on the zone one display enter in your temperature timer and press match then start. This will duplicate all of the information to zone two, Which makes it easy and super efficient. The max crisp is set at 240° and this is the only function that The temperature can't be changed, and it could be set to a maximum time of 30 minutes. The air fryer function can be set between 150 and 210° and has a time range between 1 minute to 1 hour. The roast function allows you to

~ [Read more](#)

1,486 people found this helpful

Helpful

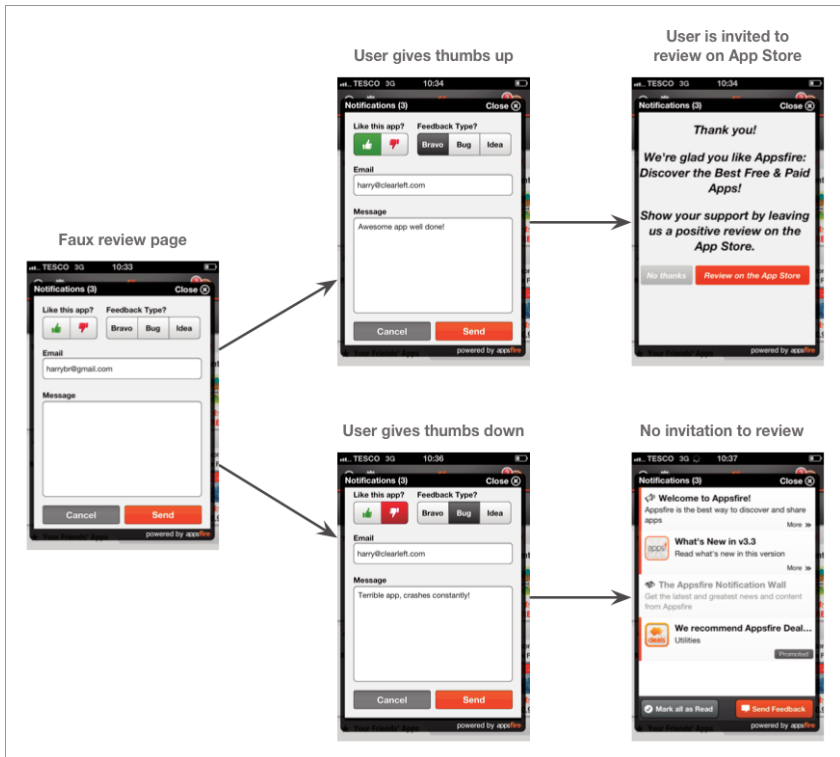
 | [Report abuse](#)

Screenshot of a review on Amazon UK, featuring the label ‘VINE VOICE’

Next to the reviewer’s name is the label ‘VINE VOICE’. The user cannot click the label or hover over the label to reveal more information – and it’s not explained on the page. If the user searches for ‘vine voice’ in the product search box at the top of the page, nothing relevant appears in the search results. Buried deep in the Amazon UK website is a ‘help library’. From there, the user can search for ‘vine voice’ and find an explanation: that reviews with this label are paid reviews,

because the reviewers were given the products for free. This is quite evidently not an adequate disclosure.

There are other ways that social proof can be manipulated. In the early days of the mobile app stores, a company called Appfire pioneered a clever approach in a product for app developers called AppBooster.¹⁶ It involved showing users a 'fake' review page in which a rating and review were requested. If users gave a thumbs up with their review, they were asked to submit it to the App Store. If users gave a thumbs down their review was transferred into an email support thread hidden away from the public – although none of this was explained to the user. You can see the steps below.



A walkthrough of the AppBooster user experience

As you can see, AppBooster was dishonest about the true purpose of the ‘thumbs up’ and ‘thumbs down’ buttons. A more honest approach would be to let users decide for themselves whether they want to leave a public App Store review or email the developer privately.

Today, this sort of manipulative technique is forbidden in the Apple and Google app stores, so it’s not seen so often. Other approaches to manipulating social proof include delaying the publication of negative reviews (holding them in a queue longer than positive reviews), or simply showing them less prominently.

SCARCITY EFFECT

Scarcity is a cognitive bias that describes the tendency for people to place greater value on resources they believe to be in limited supply. It typically influences decision-making by increasing impulsiveness and risk-taking, as people feel a sense of urgency to acquire the resource before it runs out.

One of the first and most famous studies on scarcity involves cookies – the delicious baked kind, not browser cookies. In 1975, researchers Worschel, Lee and Adwole recruited 146 undergraduate students and carried out a series of experiments.¹⁷ Participants were shown a jar of either ten cookies or two cookies, and were asked to rate how much they wanted to eat them. The results showed that participants in the two-cookie condition rated the cookies as *more desirable* and *more attractive* compared to those in the ten-cookie condition.

Then, to make matters more exciting, the researchers engaged in some theatrics during the experiment. An actor entered the room with another jar of either two or ten cookies. The actor explained that they needed to swap their jar with the one the participant was already looking at. This served to draw attention to the difference in the number of cookies, before and after. In the conditions where the number of cookies was reduced, participants rated those cookies as *even more attractive*. This just goes to show that scarcity is effective, and the effectiveness is intensified when a person’s attention is drawn to the scarcity.

In the real world, scarcity is a fact of life, and it can be very helpful to provide scarcity information to users. For example, if a user has specific dates they need to take as annual leave, it is important for them to know if their desired travel tickets are close to selling out; if they are, they'd better book them immediately or they'll miss their chance.

While honest and true messages are entirely acceptable, the scarcity effect is so powerful that it leads businesses to create fake scarcity, or to manipulate the concept of scarcity using ambiguous language, categories and user interfaces. We'll look into this further in part 3 of the book, on types of deceptive pattern.

SUNK COST FALLACY

The sunk cost fallacy is a phenomenon where individuals continue investing resources in an endeavour simply because they have already invested a significant amount in it. Even when continuing on the same path is irrational, people find it hard to let go of the resources already invested.

Research conducted by Arkes and Blumer in 1985 showed that individuals are more likely to persist in a task when they have invested resources in it, even if the investment is irretrievable and continuing the task is not rational.¹⁸ In one experiment, they gave 61 participants the following scenario. Before reading beyond the excerpt below, consider how you'd respond to this scenario.

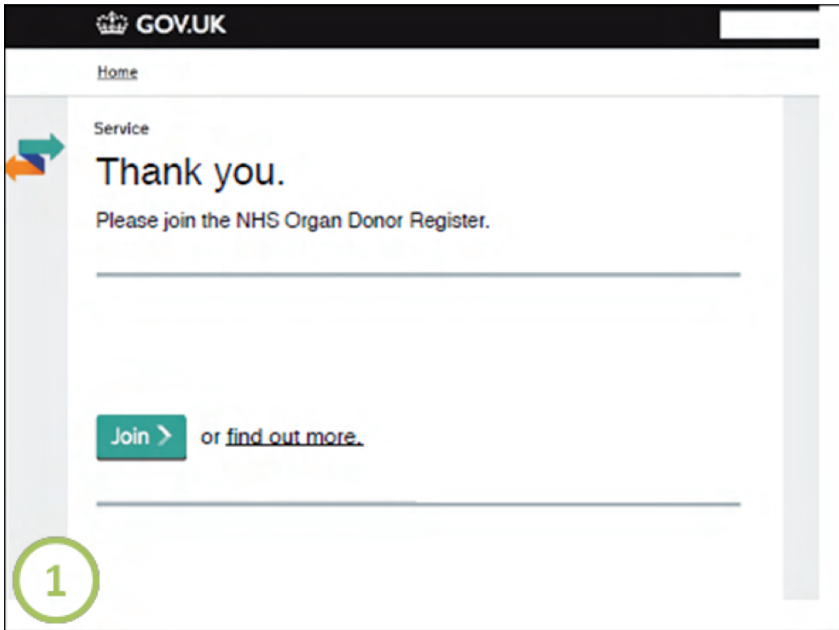
Assume that you have spent \$100 on a ticket for a weekend ski trip to Michigan. Several weeks later you buy a \$50 ticket for a weekend ski trip to Wisconsin. You think you will enjoy the Wisconsin ski trip more than the Michigan ski trip. As you are putting your just-purchased Wisconsin ski trip ticket in your wallet, you notice that the Michigan ski trip and the Wisconsin ski trip are for the same weekend! It's too late to sell either ticket, and you cannot return either one. You must use one ticket and not the other. Which ski trip will you go on?

Given the fact that all the money is now spent and cannot be retrieved, it would be irrational for you to consider the cost of the trips in making a choice. You've already worked out that you'll enjoy the Wisconsin trip so the logical choice would be Wisconsin. But did the participants in the study all pick that option? No. In fact only 46% of the respondents did. The sunk cost of the Michigan trip influenced the majority of respondents (54%).

The sunk cost fallacy is often employed in deceptive patterns by drawing users in with an attractive offer, using up their time, attention and energy going through a long-winded series of steps only to finally reveal the truth that the offer is less attractive than initially stated: the price is higher, for instance, or the terms less favourable. This will be explained further in part 3.

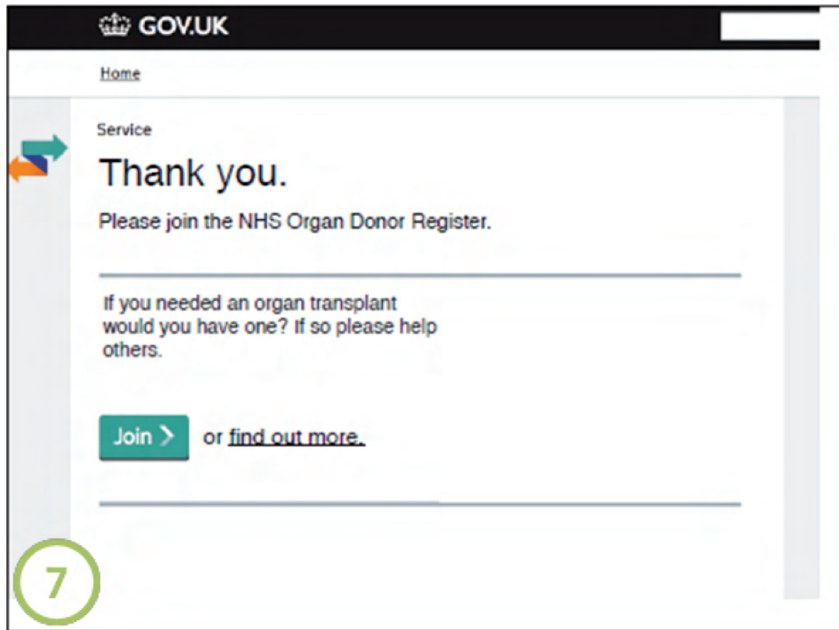
RECIPROCITY BIAS

The reciprocity cognitive bias is a phenomenon in which people tend to feel obligated to return favours to others after they have been given something. It is sometimes believed to be a form of social currency, as people may feel obligated to respond to a favour with a favour of their own. In 2013, the UK government ran a large A/B test with over 1 million website visitors, in which they tested eight different designs.¹⁹ When people had finished renewing their vehicle tax on the *gov.uk* website, they were taken a variant of this page:



A variant of the UK government vehicle tax completion page, as used in an A/B test.

They tested eight different variants of this page. The one you can see above is the control (1) and the most effective variant (7) is shown below. The two pages are identical, apart from the message in the version below: 'If you needed an organ transplant would you have one? If so please help others.'



Another variant of the UK government vehicle tax completion page, featuring a persuasive element regarding the NHS Organ Donor Register.

You might expect the effect to be small, because the text looks so unremarkable – but you’d be wrong. With the first design (1), 2.3% of people went on to register as organ donors. With the second design (7), 3.2% went on to register as organ donors. That’s one percentage point higher – or to put it another way, *one-third bigger* than the control condition.

In its report, the BIT (the UK government’s Behavioural Insights Team) refer to this design as tapping into the ‘reciprocity’ bias, a human tendency to return favours and pay back debts.²⁰ In this case, it is applied in an honest and transparent manner, but it would be deceptive if it were based on lies or misleading statements, and it’s easy to imagine it being used for nefarious ends.

PART THREE

THANKS FOR READING THE

FREE SAMPLE

If you enjoyed this, please consider buying the book, due for release on August 1st, 2023.

Available on Paperback, Kindle and DRM-free eBook formats. Visit the website for more details:

<https://www.deceptive.design/book>

ABOUT THE AUTHOR



Since 2010, Harry Brignull has dedicated his career to understanding and exposing the manipulative and deceptive techniques that are employed to exploit users online. He is credited with coining a number of the terms that are now popularly used in this research area, and is the founder of the website *deceptive.design* (formerly *darkpatterns.org*). He has worked as an expert witness on a number of cases about decep-

tive patterns, including Nichols v. Noom Inc. (case 1:20-cv-03677), Arena v. Intuit Inc. (Case 3:19-cv-02546) and FTC v. Publishers Clearing House LLC (Case 2:23-cv-04735). Harry is also an accomplished user experience practitioner, having worked for organisations that include Smart Pension, Spotify, Pearson, HMRC, and the Telegraph newspaper.

ENDNOTES

PROLOGUE

1. C-SPAN. (2021, March 25). House Hearing on Combating Online Misinformation and Disinformation [Video]. C-SPAN. <https://www.c-span.org/video/?510053-1/house-hearing-combating-online-misinformation-disinformation&live=#>

1. INTRODUCTION

1. Under advice from the Tech Policy Design Lab of the World Wide Web Foundation, I have stopped using the term ‘dark pattern’ and now use ‘deceptive pattern’. The change reflects a commitment to avoiding language that might inadvertently carry negative associations or reinforce harmful stereotypes. In this book, the term ‘dark pattern’ is used only when referring to laws, quotations and research papers that use the term.
2. This book is not a legal textbook. When the word ‘deceptive’ is used in this book, it is not intended to confer any sort of legal category or judgement. Please consider it to be intended as a like-for-like replacement of the legacy term ‘dark pattern’. In this book, the term ‘deceptive pattern’ is generally intended as a shorthand for the term ‘deceptive or manipulative pattern’.
3. Brignull, H. (2010, October 3). Dark patterns. Retrieved 3 May 2023 from <https://old.deceptive.design/> A historical snapshot of darkpatterns.org, which was recently renamed to deceptive.design
4. Flights and airline FAQs | Gatwick Airport. (n.d.). <https://www.gatwickairport.com/faqs/flights-and-airlines/>
5. Santos, D. (2018, October 9). Customer Paths and Retail Store Layout — Part 3. Aislelabs. <https://www.aislelabs.com/blog/2018/09/26/customer-paths-and-retail-store-layout-part-3>
6. Image source for figure: Gatwick Airport South Terminal Passenger Maps. (2019, December). Retrieved 3 May 2023 from <https://www.gatwickairport.com/globalassets/passenger-facilities/airport-maps/dec-2019/gatwick-airport-south-terminal-passenger-maps---dec-2019.pdf>
7. Gatwick key facts | Gatwick Airport. (n.d.). <https://www.gatwickairport.com/business-community/about-gatwick/company-information/gatwick-key-facts/>
8. Image source for figure: Brignull, H. (2010, September 28). Trick questions - dark patterns. From https://old.deceptive.design/trick_questions/ A historical snapshot of darkpatterns.org.
9. Article 4 of GDPR states “‘consent’ of the data subject means any freely given, specific, informed and unambiguous indication of the data subject’s wishes by which he or she, by a statement or by a clear affirmative action, signifies agreement to the processing of personal data relating to him or her.’”

70 ENDNOTES

10. European Parliament and Council. (2016, May 27). Regulation (EU) 2016/679. EUR-Lex. Retrieved 5 August 2022 from <https://eur-lex.europa.eu/eli/reg/2016/679/oj>.
11. Alexander, C., Ishikawa, S., & Silverstein, M. (1977). *A pattern language: towns, buildings, construction*. New York: Oxford University Press.
12. Regulation (EU) 2022/1925 of the European Parliament and of the Council of 14 September 2022 on contestable and fair markets in the digital sector and amending Directives (EU) 2019/1937 and (EU) 2020/1828 (Digital Markets Act) (Text with EEA relevance). (2022, October 12). EUR-Lex. Retrieved 5 March 2023 from <https://eur-lex.europa.eu/eli/reg/2022/1925>.
13. Regulation (EU) 2022/2065 of the European Parliament and of the Council of 19 October 2022 on a Single Market For Digital Services and amending Directive 2000/31/EC (Digital Services Act) (Text with EEA relevance). (2022, October 27). EUR-Lex. Retrieved 5 March 2023 from <https://eur-lex.europa.eu/eli/reg/2022/2065>.
14. Proposal for a Regulation of the European Parliament and of the Council on Harmonised Rules on Fair Access to and Use of Data (Data Act). (2022, February 23). European Commission. <https://eur-lex.europa.eu/legal-content/EN/TXT/HTML/?uri=CELEX:52022PC0068>
15. The California Consumer Privacy Act of 2018. (2023, January 20). State of California - Department of Justice - Office of the Attorney General. Retrieved 7 February 2023 from <https://oag.ca.gov/privacy/ccpa>.
16. Colorado Privacy Act. (2021, July 7). https://leg.colorado.gov/sites/default/files/2021a_190_signed.pdf
17. By *enforcer* I mean any entity that acts to ensure compliance with legal regulations and protects consumers from deceptive patterns, either directly or indirectly. Many regulators are enforcers (e.g. the Federal Trade Commission, the Competition and Markets Authority), but enforcement also can occur via private law firms, consumer advocacy groups, and others.

2. A PRIMER ON DESIGN INDUSTRY TERMINOLOGY

1. A number of the patterns described in this book are not deceptive under the FTC Act's definition of the term *deceptive*, e.g. confirm-shaming, nagging or forced action. Those sorts of patterns are better described as manipulative. Since this book is not a legal text, I have stuck with the term *deceptive patterns* throughout and intend it as a synonym to the term *dark patterns* as used by the FTC and other parties in the US.

3. THE RISE OF DECEPTIVE PATTERNS

1. Stanford Digital Civil Society Lab. (n.d.). Dark Pattern Tipline. Retrieved 3 August 2022 from <https://darkpatternstipline.org/>
2. Stevens, M. (2016). *Cheats and deceptions: How animals and plants exploit and mislead*. Oxford University Press.
3. Underhill, P. (1999). *Why we buy: The science of shopping*. Simon & Schuster.
4. You may also have heard of split testing and multivariate testing (MVT). Both are conceptually similar to A/B testing with some technical differences.
5. Hopkins, Claude C. (1923) Scientific advertising. <http://www.scientificadvertising.com/ScientificAdvertising.pdf>

4. FROM HOMO ECONOMICUS TO HOMO MANIPULABLE

1. Simon, H. A. (1986). Rationality in psychology and economics. *The Journal of Business*, 59(4), S209–S224. <http://www.jstor.org/stable/2352757>
2. Nobel Prize in Economic Sciences 2017: <https://www.nobelprize.org/prizes/economic-sciences/2017/press-release/>
3. Thaler, R. H., & Sunstein, C. R. (2008). *Nudge: Improving decisions about health, wealth, and happiness*. Yale University Press.
4. Wickens, C.D., Gordon, S., & Liu, Y. (1997) *An introduction to human factors engineering*. Longman. https://openlibrary.org/works/OL2728752W/An_introduction_to_human_factors_engineering
5. Jarovsky, L. (2022, March 1). Dark patterns in personal data collection: Definition, taxonomy and lawfulness. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4048582

II. EXPLOITATIVE STRATEGIES

1. Gray, C. M., Kou, Y., Battles, B., Hoggatt, J., & Toombs, A. L. (2018). The dark (patterns) side of ux design. *Proceedings of the 2018 CHI conference on human factors in computing systems*. <https://doi.org/10.1145/3173574.3174108>

5. EXPLOITING PERCEPTUAL VULNERABILITIES

1. Purves, D. (2001). *Neuroscience*. Palgrave Macmillan.
2. Gleitman, H., Gross, J., & Reisberg, D. (2011). *Psychology*. WW Norton & Company.
3. Lime hawk-moth | Cumbria Wildlife Trust. (n.d.). <https://www.cumbriawildlifetrust.org.uk/wildlife-explorer/invertebrates/moths/lime-hawk-moth>
4. Image source for figure: Sale, B. (2018). Lime hawk-moth (*Mimas tiliae*). Flickr. https://flickr.com/photos/33398884@N03/40578533840_cc-by-2.0
5. W3C. (n.d.). G17: Ensuring that a contrast ratio of at least 7:1 exists between text (and images of text) and background behind the text | Techniques for WCAG 2.0. w3.org. Retrieved 3 August 2022 from <https://www.w3.org/TR/WCAG20-TECHS/G17.html#G17-tests>
6. WebAIM. (n.d.). WebAIM: Contrast checker. [webaim.org](https://webaim.org/resources/contrastchecker/). Retrieved 3 August 2022 from <https://webaim.org/resources/contrastchecker/>
7. Atrash, D. (2022, February 8). Understanding web accessibility standards: ADA, Section 508, and WCAG compliance. Medium. <https://bootcamp.uxdesign.cc/understanding-web-accessibility-standards-ada-section-508-and-wcag-compliance-143cfb8b691e>
8. *Arena v. Intuit Inc.* Case No. 19-cv-02546-CRB. (2020, March 12). Casetext. Retrieved June 29, 2023, from <https://casetext.com/case/arena-v-intuit-inc>
9. *Arena v. Intuit Inc.* Case No. 19-cv-02546-CRB. (2020, March 12). Casetext. Retrieved June 29, 2023, from <https://casetext.com/case/arena-v-intuit-inc>
10. When taken to the Ninth Circuit, this decision was actually reversed, which demonstrates some ambiguity in the nature of the case and US law.
11. Nouwens, M., Liccardi, I., Veale, M., Karger, D., & Kagal, L. (2020). Dark patterns after the GDPR: Scraping consent pop-ups and demonstrating their influence.

- Proceedings of the 2020 CHI conference on human factors in computing systems. <https://doi.org/10.1145/3313831.3376321>
12. The Behavioural Insights Team. (2014a, April 11). EAST: Four simple ways to apply behavioural insights. BIteam.com, 11th Apr 2014. Retrieved June 17, 2023, from http://www.bi.team/wp-content/uploads/2015/07/BI-Publication-EAST-FA_WEB.pdf
 13. Image source for figure: How can a letter encourage us to pay our parking fines? (4 March 2016). The Behavioural Insights Team. Retrieved 17 October 2022 from <https://www.bi.team/blogs/how-can-a-letter-encourage-us-to-pay-our-parking-fines/>
 14. In the plain letter condition, 14.7% of recipients made a payment; while in the red stamp condition, 17.8% of recipients made a payment. The total sample size was 48,445. BIT did not report the numbers assigned to each condition. http://www.bi.team/wp-content/uploads/2015/07/BI-Publication-EAST-FA_WEB.pdf
 15. Williams, R. (2015). *The non-designer's design book: Design and typographic principles for the visual novice*. Amsterdam University Press.

6. EXPLOITING VULNERABILITIES IN COMPREHENSION

1. PIAAC. (n.d.). The Programme for the International Assessment of Adult Competencies. Retrieved January 24, 2023, from <https://nces.ed.gov/surveys/piaac/about.asp>
2. Infographics. (n.d.). PIAAC Gateway. Retrieved 24 January 2023 from <https://www.piaacgateway.com/infographics>
3. Image source for figure: Justin Hurwitz, Americans at risk: Manipulation and deception in the digital age. (Written testimony of Justin Hurwitz) (2020) <https://www.congress.gov/event/116th-congress/house-event/LC67008/text?loclr=cga-committee>
4. Krug, S. (2006). *Don't make me think! A common sense approach to web usability*. New Riders.
5. This applies to websites and apps in which the user's goals are at odds with the volume of content ('I want to find a way to quickly get through this content so I complete my task'). The obvious exception here is novels and long-form content, when the user's goal is to study or enjoy every single word, despite the high cost of doing so in time, attention and energy.
6. Morkes, J., & Nielsen, J. (1997, January 1) Concise, SCANNABLE, and objective: How to write for the web <https://www.nngroup.com/articles/concise-scannable-and-objective-how-to-write-for-the-web/>
7. Nielsen, J. (1997, September 30). How users read on the web <https://www.nngroup.com/articles/how-users-read-on-the-web/>
8. Pernice, K., Whitenon, K., & Nielsen, J. (2014). How people read online: The eyetracking evidence <https://www.nngroup.com/reports/how-people-read-web-eyetracking-evidence/>
9. Pirolli, P., & Card, S.K. (1999). Information foraging. *Psychological Review*, 106(4), 643–675. <https://doi.org/10.1037/0033-295X.106.4.643>
10. Federal Trade Commission. (2022, September 15). Bringing dark patterns to light - FTC staff report. Retrieved 1 January 2023 from <https://www.ftc.gov/reports/bringing-dark-patterns-light>

11. Luguri, J., & Strahilevitz, L.J. (2021, January 1). Shining a light on dark patterns. *Journal of Legal Analysis*, 13(1), 43–109. <https://academic.oup.com/jla/article/13/1/43/6180579>

7. EXPLOITING VULNERABILITIES IN DECISION-MAKING

1. Society for Judgment and Decision Making. (n.d.). Retrieved 23 January 2023 from <https://sjdm.org/>
2. Wylie, C. (2020). *Mindf*ck: Cambridge Analytica and the plot to break America*. Penguin Random House.
3. Ariely, D. (2010). *Predictably irrational: The hidden forces that shape our decisions*. Revised and expanded edition. Harper Perennial.
4. Sloman, A. (1989). Preface. In M. Sharples, D. Hogg, S. Torrance, D. Young, & C. Hutchinson, *Computers and thought: A practical introduction to artificial intelligence*. Bradford Books. <https://www.cs.bham.ac.uk/research/projects/cogaff/personal-ai-sloman-1988.html>
5. Benson, B. (2016, September 1). Cognitive bias cheat sheet: An organized list of cognitive biases because thinking is hard. *Better Humans*. Medium. Retrieved 23 September 2022 from <https://betterhumans.pub/cognitive-bias-cheat-sheet-55a472476b18>
6. Cialdini, R.B. (2001). *Influence: Science and practice*. Allyn and Bacon. The book details '7 weapons of influence': scarcity, authority, social proof, sympathy, reciprocity, consistency and unity.
7. Schüll, N.D. (2014). *Addiction by design: Machine gambling in Las Vegas*. Amsterdam University Press.
8. 250 best A/B testing ideas based on neuromarketing. (n.d.). Convertize.com. Retrieved 31 January 2023 from <https://tactics.convertize.com/principles>
9. Thaler, R.H. (2015). *Misbehaving: The making of behavioural economics*. Penguin Books Ltd.
10. Servicio Nacional del Consumidor [SERNAC]. (2022, March). Policy paper on cookies consent requests: Experimental evidence of privacy by default and dark patterns on consumer privacy decision making. Retrieved 28 January 2023 from https://icpen.org/sites/default/files/2022-05/SERNAC_Policy_Paper_Cookies_Experiment.pdf
11. Tversky, A., & Kahneman, D. (1974). Judgement under uncertainty: Heuristics and biases. *Science*, 185, 1124–1131. <https://doi.org/10.1126/science.185.4157.1124>
12. Tversky, A., & Kahneman, D. (1981). The framing of decisions and the psychology of choice. *Science*, 211, 453–458. <https://doi.org/10.1126/science.7455683>
13. Ariely, D. (2010). *Predictably irrational: The hidden forces that shape our decisions*. Revised and expanded edition. Harper Perennial.
14. Hallsworth, M., List, J.A., Metcalfe, R.D., & Vlaev, I. (2017). The behavioralist as tax collector: Using natural field experiments to enhance tax compliance. *Journal of Public Economics*, 148, 14–31. <https://doi.org/10.1016/j.jpubeco.2017.02.003>
15. Ninja Foodi Air Fryer. (n.d.). Amazon.co.uk. Retrieved 4 February 2023 from <https://www.amazon.co.uk/Ninja-Foodi-Fryer-Dual-Zone/dp/B08CN3G4N9/>
16. Brignull, H. (2021, May 21). Manipulating app store reviews with dark patterns. 90 Percent of Everything. Retrieved 4 February 2023 from <https://90percentofeverything.com/2012/05/21/manipulating-app-store-reviews-with-dark-patterns/>

17. Worchel, S., Lee, J. W., & Adewole, A. (1975). Effects of supply and demand on ratings of object value. *Journal of Personality and Social Psychology*, 32(5), 906–914. <https://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.822.9487>
18. Arkes, H.R., & Blumer, C. (1985). The psychology of sunk cost. *Organizational Behavior and Human Decision Processes*, 35(1), 124–140. [https://doi.org/10.1016/0749-5978\(85\)90049-4](https://doi.org/10.1016/0749-5978(85)90049-4)
19. Behavioural Insights Team with Cabinet Office, Department of Health, Driver and Vehicle Licensing Agency, & NHS Blood and Transplant. (2013, December 23). Applying behavioural insights to organ donation. Behavioural Insights Team. Retrieved 17 October 2022 from <https://www.bi.team/publications/applying-behavioural-insights-to-organ-donation/>
20. For more information on the use of cognitive biases in persuasion, read Cialdini, R.B. (2001). *Influence: Science and Practice*. Allyn and Bacon. The book details ‘7 weapons of influence’: scarcity, authority, social proof, sympathy, reciprocity, consistency and unity.