The impact of enhanced product imagery on an e-commerce business

Gonçalo Barbosa Camiña Barata Rebelo

Dissertação de Mestrado

Orientador na FEUP: Prof. Mário Amorim Lopes

U. PORTO
FEUP FACULDADE DE ENGENHARIA
UNIVERSIDADE DO PORTO

Mestrado Integrado em Engenharia e Gestão Industrial

2019-06-25

"Luxury is the ease of a t-shirt in a very expensive dress." - **Karl Lagerfeld**

The impact of enhanced product imagery on an e-commerce business

Abstract

Any business in the fashion retail sector needs to constantly adapt to the market. Trends are set overnight, which means that companies must be very agile in terms of acquiring the most sought after products. Furthermore, since the level of competition in the business is extremely high – a lot of companies selling the very same products – there is a high need for differentiation. One of the greatest methods for staying competitive is providing photographic content of the highest quality.

Since images are so unique and complex, it is sometimes challenging to understand how they can be improved. To achieve this, live testing may be of immense help. By assessing the impact of the variation of certain image factors on some of the businesses main performance metrics, companies can more easily identify valuable opportunities.

The objective of this thesis is not only to study a particular image variable – number of shots – but also to understand in detail the testing process and the benefits associated with it. For this purpose, an A/B test was executed and its results were analysed, leading to a positive and, to a certain extent, surprising conclusion – by providing extra shots of Coats, Denim, Dresses and Trousers more visitors actually purchased the products.

Resumo

As empresas no ramo do retalho de moda necessitam de se adaptar constantemente ao mercado. Modas formam-se da noite para o dia, pelo que as empresas precisam de ser muito ágeis no que toca à aquisição dos produtos mais procurados. Para além disso, uma vez que o nível de competição é extremamente elevado – muitas empresas a vender exatamente os mesmos produtos –, há uma grande necessidade de diferenciação. Uma das melhores formas de se obter diferenciação é disponibilizar conteúdo fotográfico da melhor qualidade.

Como as imagens são tão unicas e complexas é, por vezes, um desafio perceber como é podem ser melhoradas. Para este fim, os testes em direto são uma ferramenta útil. Ao avaliar o impacto da variação de certos fatores de imagem em algumas das principais métricas de performance do negócio, as empresas conseguem, mais facilmente, identificar oportunidades valiosas.

O objetivo desta tese é não só estudar uma variável de imagem em particular – número de fotografias –, mas também perceber em concreto o processo de teste e os benefícios do mesmo. Para este efeito, um teste A/B foi executado e os seus resultados analisados, levando a uma conclusão positiva e, de certa forma, surpreendente – ao disponibilizar fotografias extra de Calças, Casacos, Ganga e Vestidos mais visitantes efetivamente compraram os produtos.

Acknowledgments

In the first place, I would like to start by thanking every colleague in Farfetch who has, in any possible way, contributed to the success of this project – especially Frederico Alves, Sofia Leal and Tiago Rocha, to whom I owe the opportunity to be involved in the project and much of its success.

Furthermore, I would like to express my gratitude towards Prof. Mário Amorim Lopes whose availability and constant guidance ensured I kept following the right path.

Finally, I want to thank all my family and friends, who have supported me throughout the development of the project.

Table of Contents

List of Acronyms

UAE – United Arab Emirates

US – United States

CAGR – Compound Annual Growth Rate

PLP – Product Listing page

PDP – Product Display page

QC – Quality Control

A2WR – Add to Wishlist Rate

A2BR – Add to Bag Rate

SCR – Started Checkout Rate

CR – Conversion Rate

RR – Return Rate

SE – Standard Error

LB – Lower Bound

UB – Upper Bound

ROI – Return On Investment

GMV – Global Merchandised Value

Table of Figures

1 Introduction

In this chapter, the context and motivation of this thesis are introduced, including a brief description of Farfetch and its departments involved in this experience. A very high-level description of the document's structure is outlined at the end of the chapter.

1.1 Context and Motivation

The luxury fashion market is growing at an incredible rate. According to a forecast by (Bain & Company, 2019), this market is expected to grow by 60 to 100 billion dollars until 2025 (Figure 1).

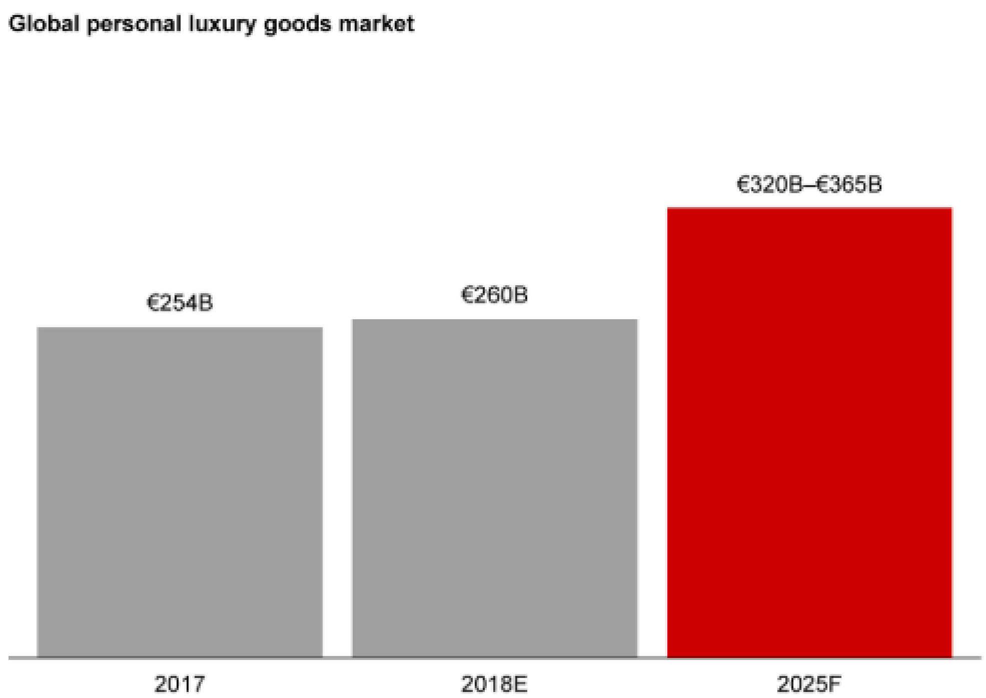

Figure 1 - Global personal luxury goods market volume

Currently, online sales represent around 10% of this market, totaling 26 billion dollars. By 2025, however, the online slice is expected to increase to 25% which by then, taking into account the total market's growth to 320 to 365 billion dollars, would consist of 80 to 90 billion dollars (Bain & Company, 2019a). Summing up, this growth could represent an increment of 55 to 65 billion dollars in the sector Farfetch competes at which is – obviously - an incredible opportunity.

But it is not enough to have the opportunity. The fact that the market is growing also means that it is very likely that competitors try to expand, and new players, guided by the attractiveness of the opportunity, join the market.

Companies in the sector need to attract new customers in all the ways possible, while also retaining their existing ones. In environments of high competitiveness this might prove a very difficult task, thus the need for differentiation at every level.

Even more crucial than having a high user traffic (potential customers) on the website, is having a converting customer base. In Farfetch's case, customers need to enter the website and feel inspired to purchase. And one of the best ways of convincing someone to purchase on a website is having high-quality product imagery.

Imagery is, undoubtedly, a key element in e-commerce. It is understood that better imagery drives more sales (Cardew, n.d.), as customers feel more engaged in the purchase activity and have a higher perception of the products, quality-wise. This is especially true in the luxury fashion sector, where there is usually a lot of reluctancy to buy, mainly due to the fear of misunderstanding how the product really looks and feels – something that can be partly solved by providing a better representation of the product.

However, on a business with the volume of Farfetch's it is not advisable to implement any kind of major change without first predicting or – preferably - testing its impact. This is especially valid for image-level changes, as they are one of the main means of communication between the company and its customers.

The motivation for this study is, therefore, understanding how changes in imagery can impact the business, through advanced analytics.

1.2 Farfetch Overview

Farfetch is an online luxury fashion retailer that sells products from over 1000 partners worldwide.

Founded in 2007 by José Neves, the platform has now offices in 7 countries (Brazil, China, England, Japan, Portugal, UAE and US), with its headquarters located in London.

In September 2018, Farfetch was listed on the New York Stock Exchange, being valued at over $5.8 billion.

Besides the main business unit - Farfetch.com -, the platform also includes the boutique and e-retailer Browns Fashion, as well as the Black & White service, which provides a "ready-to-use" e-commerce solution to brands and stores. The latest Black & White partner, and one of the biggest, is the English store "Harrods".

Currently, Farfetch has four production centers worldwide. These are located in Guimarães (Portugal), Los Angeles (USA), São Paulo (Brazil) and Wong Chuk Hang (Hong Kong). Around 80% of Farfetch's production occurs in Guimarães, while the remainder is allocated to the other three offices. The production guidelines are, nevertheless, aligned between all locations.

The Operations division within Farfetch is divided into 7 distinct teams: Back-Office Product handles the infrastructure that supports the platform's functioning; Creative Operations is responsible for the creation of written and photographic content for all functional areas within the FF group; Customer Excellence guarantees the satisfaction of customers; Logistics is responsible for planning, implementing and controlling the effective flow of all services; Office Operations manages the functioning of all offices; Operations Strategy supports all the areas of the business via several projects; and finally, Platform Operations coordinates the Black & White business unit.

Creative Operations and Operations Strategy partnered up, and with the support of several other teams (inside and outside Operations) started this imagery enhancement project, with the goal of improving the quality of the photographic content presented to the website's visitors. This project consists of multiple initiatives, targeting several image-related aspects, such as the alignment of products on the listing pages and the number of product details shown.

1.3 Objectives

The objective of this dissertation is to test the implementation of extra shots of product details (describing the implementation process from start to finish), understand how they influence the customers' experience, determine how can their impact on certain performance metrics be assessed, and finally, proceed to evaluating its impact.

1.4 Organization

Following this chapter, the document is structured as follows.

Chapter 2 features the background in which this thesis is inserted. Initially, the luxury sector's landscape is depicted, with some information and statistical facts about the luxury consumer and market being provided. Afterwards, we provide an overview of how e-commerce functions, the emotional behaviour of consumers, and a brief account of fashion e-commerce imagery – where the best practices are described. Finally, we analyse the importance of testing.

Chapter 3 contextualizes the problem. It includes an enumeration of imagery variables, as well as a brief description of Farfetch's production process.

Chapter 4 describes the methodology adopted to implement the improvement. It starts with the choice of the variable to be used as the object of study, and proceeds to describe the data gathering methods, as well as the formulas used to assess the relevant performance metrics. The monitoring framework is also detailed here.

Chapter 5 consists of the results obtained from the analysis described in Chapter 4. It also includes the decision-making process regarding the progression from testing to rollout.

Chapter 6 showcases the conclusions drawn from the study. Furthermore, a comparison – imagery-wise – is established with the closest competitors, from which Farfetch's weakest points are determined and considered as next-in-line to be improved.

2 Background

2.1 The Luxury Sector Landscape

This subtopic is dedicated to the description of the current state of the luxury sector, as well as the prediction of its trends in the following years. The analysis is divided into luxury consumer and luxury market.

2.1.1 The Luxury Consumer

The luxury consumer base is shifting. Every year, the percentage of luxury consumers is transitioning from generation. For instance, while in 2017 Generation X consumers constituted around 15% of the luxury segment, in 2018 that number was reduced to zero (BAIN & COMPANY Media Center, 2018).

It is a known fact that luxury consumers' behaviour differs greatly from that of regular ones. Although trends affect both types of consumers, the latter is driven mostly by necessity, while the former is very influenced by uniqueness (Jain, Khan, & Mishra, 2017).

Luxury products are characterized by their reduced supply (Janssen, Vanhamme, Lindgreen, & Lefebvre, 2014; Y. Kim, 2018), and therefore are very attractive to individuals who possess a high need for uniqueness (Lynn 1991), such as the luxury consumer.

According to (Snyder, C.R. and Fromkin, 1980) uniqueness theory, people find both high levels of similarity and dissimilarity unpleasant, therefore seeking, alternatively, to be moderately distinct from others. In practical terms, this means that as people start feeling similarities to others, they become motivated to distinguish themselves, thus establishing their dissimilarity and uniqueness (Lynn & Harris, 1997).

The impact of social media and all forms of information available online on the luxury consumer's journey has been constantly growing. According to (BCG & Altagamma, 2018), social media is the primary source of information of the modern luxury consumer, and thus, the channel of primary impact. Currently, Facebook, Instagram, WeChat, Weibo and QQ are on top of the social media charts. A study by (McKinsey & Company, 2016) concluded that around 80% of personal luxury goods sales were influenced by online (Figure 2).

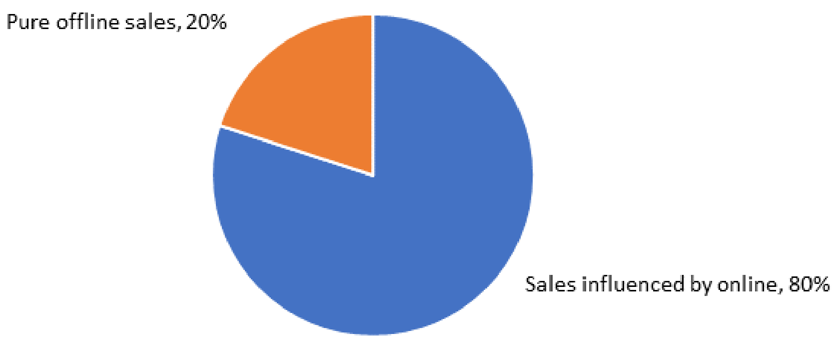

Figure 2 - Percentage of sales that are influenced by the "online" factor

2.1.2 The Luxury Market

In 2018, the total personal luxury goods market grew at a Compound Annual Growth Rate (CAGR) of 5% - achieving an astounding 1.3 trillion dollars - and was forecast to keep growing at a “new normal” rate of 4% - 5% until 2020.

Sales of luxury cars keep dominating the market, displaying a growth of 5% (considering constant rates) towards 552 billion dollars. Luxury experiences remained very sought by customers – 5% growth on luxury hospitality, 6% growth on gourmet food/fine dining and 7% growth on luxury cruises.

Worldwide, the Americas and Europe dominated the personal luxury goods market with the respective shares of 32% and 33%. Despite this fact, the spotlight is on China that is the key growth driver, featuring a share of 8%. This growth is mainly explained by the rise of the middle class, as well as the strong government initiatives to encourage local consumption (Bain & Company, 2019b).

Online luxury fashion – the main market where Farfetch competes – grew nearly 22% to 30 billion dollars. It now represents around 10% of all luxury sales. The Americas comprise 44% of online sales.

2.1.3 Trends that will shape the Luxury Market until 2025

According to (Bain & Company, 2019c), certain trends are set to have a great impact on the luxury market until 2025.

The first prediction made by Bain is the looming of the Chinese consumer. As seen in 2018, the Chinese are growing at a rapid rate on the luxury market. In fact, Bain predicts that by 2025 they will account for 46% of the global market – as opposed to the current 33% they now hold.

Furthermore, the magnitude of “Digital”, which is already considerable, is going to ramp up in the next years. By 2025, Bain expects the online channel to grow to 25% of all luxury sales. It also predicted that 50% of all luxury purchases will be digitally enabled and nearly 100% of them will be somehow influenced by online interactions.

Young consumers are also set to dominate the market. In fact, according to Bain, Generations Y and Z together will comprise approximately 55% of the luxury market in 2025 – contributing an outstanding 130% to the market growth in the upcoming years.

2.2 E-commerce overview

The term "commerce" only describes buying and selling transactions between business partners. If this definition of commerce were used, the term "electronic commerce" would be too narrow. Thus, the need for a term such as "e-business" instead (Turban et al., 2018).

E-business presents a broader description of e-commerce - not only purchasing and retailing but also conducting other kinds of businesses online, such as servicing customers, delivering e-learning, collaborating with business partners and conducting electronic transactions within organizations (Turban et al., 2018a).

2.2.1 E-commerce and Consumer Emotions

E-commerce and consumer emotions are two intimately related concepts. Ultimately, the success of an e-commerce business lies totally on the satisfaction of their customers. The client's emotions throughout the whole journey of an online purchase have a very important effect on key performance metrics, such as conversion and customer retention rates.

Regarding this topic, Lievonen's publication "Consumer Emotions and E-commerce" (2017) is analysed. The author divides his analysis of consumer emotions into four topics: website design effects on consumer emotions; pre-consumption emotions; consumer emotions during actual online shopping encounters; and finally, post-purchase emotions.

According to the author, emotions play a very important role in what is defined as the "perception" stage of websites, having an impact on the consumers' attitude toward the marketplace, as well as on their level of involvement and purchase intention. Certain studies suggest that the use of human images instead of generic ones, as well as music and colours, stimulate higher levels of satisfaction (Bui & Kemp, 2013; Cheng, Wu, & Yen, 2009; Ding & Lin, 2012; Wang, Yang, Wang, & Ma, 2014; Wu, Cheng, & Yen, 2008).

Concerning the pre-consumption phase, the author mentions the influence of the information available online (often through user reviews and social media discussions). He specifies the topic of past occurrences of trust violations on the retailers' side, and how reconciliation efforts are valued by customers when considering their alternatives (Choi & Nazareth, 2014).

Lievonen also mentions the importance of emotional and rational routes, along with the quality of the e-service itself, as key aspects of online shopping (Hsin Chang & Wang, 2011; Wang et al., 2014). Uncertainty when it comes to the checkout procedures is mentioned as a common issue for e-commerce retailers in general. Emotions have a strong impact on the consumers' final decision to whether or not provide certain required information (such as credit card numbers) (Li, Sarathy, & Zhang, 2014). In order to reduce the likelihood of consumers abandoning their initiated purchases, a very reassuring and trustworthy relationship between them and the retailer must be established (Peiris, Kulkarni, & Mawatha, 2015). In sum, when it comes to achieving positive conversion rates, the emotions felt on this phase are of critical importance.

Finally, in relation to post-purchase, the author establishes that the emotions felt after completing a purchase have a lot of weight on the consumers' decision about whether or not to return to the website and engage in further transactions. As mentioned by (Kuo & Wu, 2012), positive emotions increase post-satisfaction, which in turn affects post-purchase intentions. Once again, we can clearly see the impact of consumers' emotions on business KPIs, in this case specifically on customer retention. Returns usually have a very negative impact on the customers' emotions.

2.3 Fashion E-commerce Imagery

Although (Lievonen, 2018) presents a very interesting and complete analysis on consumer emotions, not enough focus was given to the imagery component of e-commerce, one of the components that has the greatest potential to influence customers. The author does mention it at some point but does not analyse its potential impact in detail.

"We live in a visual world. A customer walks into a brick and mortar store because they like what they see from the outside. Similarly, when customers land on your store page, their attention gravitates to your images first and, if they like what they see, they keep on browsing and, hopefully, make a purchase. That is why product photography is essential to the success of any e-commerce operation." (Cardew, n.d.).

Eye tracking studies show that store visitors are first engaged by visual elements, which, if adequate, make them more likely to stay in the site and explore it (Ciotti, 2013).

Cardew defends a few key points:

In the first place, photography and branding are strongly intertwined. When a brand's imagery resonates its values and those of their customers/followers, it adds value to their lives, making them loyal promoters. Images enrich content and speak to the target audience. A very good example of this is the fact that if one thinks about a white polar bear in a red background, it instantly recalls "Coca-Cola".

Secondly, that photography is a crucial part of any creative strategy. Images are a reliable source to improve understanding and increase engagement, and should therefore be timeless, emotional and dynamic. A favourable perception is created by ensuring that the imagery used has these characteristics, which ultimately drives sales.

The author's point of view is corroborated by some statistical facts gathered by (L. Kim, 2015): the average person gets distracted in 8 seconds (although 2.8 seconds is enough for some people); 81 percent of people only skim the content they read online; people form a first impression in a mere 50 milliseconds; posts that include images produce 650 percent higher engagement than text-only posts; posts with videos attract 3x more links than text-only posts; people are 85 percent more likely to buy a product after viewing a product video.

Another interesting statistical fact is that 65% of the human population are visual learners – this means that around 65% of consumers can be influenced through the use of imagery (Khan, 2018).

2.3.1 State of the art Product Photography

Product photography is divided into "basic" (product only) and "in-context" or "lifestyle". Depending on the type of product to be showcased, one or both types can be applied. Since this project is inserted in the fashion area, it made sense to only focus on products related to it.

When it comes to fashion photography, the basic type of photography is mandatory. In-context/lifestyle photography, being more expensive and time-consuming is treated as a complementary option, strongly recommended for high-end fashion products and/or considerable sales volumes (which is the case for Farfetch).

Since there are considerable differences between the guidelines of the aforementioned types of product photography, we analysed them separately.

Basic/Product only

When taking basic product photographs, there are a few standard elements to be taken into account.

First off, a white or very light background is established as the gold standard for product photography. It has several advantages, namely promoting consistency across the store page (as it helps to create a uniform look), as well as being the easiest format to be published by third parties, such as magazines or blogs. This type of background can be achieved by either post-production (Adobe Photoshop is the most used software for this) or by photographing the product in a physical white background - usually, a sweep (Appendix A) - resorting only to post-production in order to perfect the image (Cardew, n.d.). The latter is strongly recommended not only because it drastically reduces the need for post-production editing, but also because the white surface reflects some light onto the product, giving it a crisper look.

To what concerns lighting, there are two major groups to consider: natural and artificial light. Cost-wise, using the Sun's light to shoot is preferable, since there is no need to invest in studio lights and, consequently, electricity. Natural light photoshoots can be either indoors or outdoors. When shooting indoors, the recommended timeframe is around noon (when the light is at peak intensity), and the best way to do this is by placing the sweep/shooting table right beside a window and placing a white sheet in front of the latter (to slightly dim the light). For outdoor shots, on the other hand, very intense light should be avoided, being preferable to shoot during the golden hours (right after sunrise and right before sunset) or, if shooting around noon, doing so in a shaded area. It is very unusual to utilize white physical backgrounds when shooting outdoors. Artificial light shots, on the other hand, resort to, as the name indicates, artificial lights. The ideal setup consists of three lights – two on either side of the shooting area, and one above it – although two lights are usually enough. A basic set of lights can be purchased by around £50. More advanced producers can also use lightboxes – cubical structure filled with light – to obtain the best results illumination-wise. This type of equipment is mostly used when shooting jewelry (Cardew, n.d.).

Using natural light is recommended for beginner product photographers and/or low volume of products to shoot. When the product volume is higher, resorting to artificial lights is the most practical option, while also promoting imagery consistency among all items.

Obviously, photography is not possible without a camera. Surprisingly, the latest generation of smartphones can be used in product photography and are a great option for beginners. Nevertheless, the ideal device should be a DSLR 50mm camera, such as a Canon EOS 70D or Nikon D3200. In order to obtain consistent results and a reduced number of blurred photos, the camera should be mounted on a tripod (Cardew, n.d.).

Studio props such as mannequins should be considered as an additional option. They are a good option to showcase medium to large pieces of clothing, giving them some bulk. Still, as the mannequin itself is not very aesthetical, it should be removed from the photo in post-production. To achieve this, typically a ghost mannequin technique is used which consists in taking photos of the front and back of the product and merging them, thus removing the mannequin – which takes some skill and time (Cardew, n.d.).

In-context/Lifestyle

As previously mentioned, this type of product photography is usually reserved for top fashion brands and/or high-volume retailers.

When it comes to fashion, in-context/lifestyle photoshoots resort primarily to a model.

Model photography is more complex and expensive than basic photography. The complexity comes from the added factors, such as makeup definition and model expression and pose, while the added costs come from the model itself – models usually charge by piece photographed, with varying prices according to the degree of body exposure – as well as the need to hire stylists and makeup artists (Cardew, n.d.).

Besides the increased costs and shooting times, when done correctly and adequately, using models can also present a substantial increase in the volume of sales. Using models is the best way of telling a story about the products, and therefore, creating an emotional link with the customer – who is induced, even if subconsciously, to imagine him/herself wearing the clothes just as the model. Furthermore, having models/influencers showcase a brand's product is a proved way of boosting its image. Last but not least, showing customers how the clothes fit on a real person is a great way to encourage them to convert, especially in the luxury sector where, due to the higher price tag associated with products, they are even more reluctant to buy - fearing, for example, buying something that doesn't fit as expected (Cardew, n.d.).

Just as basic photography, model photography can be practiced either indoors or outdoors, and resort to natural or artificial lighting. When shooting indoors, the equipment might somewhat differ from that of basic product photography. Larger sweeps - possibly from ceiling to floor - might be needed to accommodate the model's full body, as well as bigger lights (to cover larger scenarios) and taller tripods (Cardew, n.d.).

Fashion e-commerce on Instagram

A sub-topic must be dedicated to Instagram e-commerce because besides being the most image-focused social platform, since its acquisition by Facebook in 2012 it includes a wide set of commercial features, of which the Buying functionality, as well as the Shopify integration, should be highlighted. Sellers using the platform are not only allowed to increase awareness, but also sales through an extra distribution channel.

Having said that, when selling on this platform it is very important to choose the adequate photos to showcase products.

The audience on this platform appreciates storytelling, which can be achieved, for example, by using series of photos. Furthermore, sellers should also focus on style over quality (resorting to an alternative, lifestyle photography instead of highly art directed one) – ensuring that the true brand's personality is very visible (consistency and minimalism) (Cardew, n.d.).

2.4 The importance of testing

In any fast-paced business there is little to no margin for error. Companies need to keep up with the trends, while also innovating and generating more demand in order to fend off the competition (Simon, 2018). Innovation is a synonym for change, but this change cannot be random – the most important (or profitable) changes must be prioritized.

Even though all companies face this challenge, it is the largest ones that have the hardest task. Filled with bright minds full of ideas, new possibilities are generated on a daily basis, which leads to an increased decision complexity. In inflexible companies, this overwhelming amount of options can be a recipe for disaster – in order not to lose time thinking which opportunities should be followed, they might choose a random path, or even no path at all – which is why it is of paramount importance to have the capability to distinguish the best opportunities (Kohavi & Thomke, 2017).

The impact of change can be predicted through forecasts, but there is really no better way to assess an opportunity than with live tests using real customer behaviour analysis.

2.4.1 A/B Testing

One of the simplest testing methods, which also returns incredibly accurate results, is designated "A/B testing". Also called "split testing", it is used to test the impact of various factors in several forms of communication. In this type of test, the experimenter designs two experiences: control version "A", which represents the current system – also designated "Champion" – and an alternative version "B", presenting a new possible scenario – the "Challenger". It is also possible to include more than 2 versions (A/B/C/… testing) (Kohavi & Thomke, 2017a).

Although its first web usage dates back to the year 2000 - when Google used it to ascertain the optimum number of search results displayed by page – we have been using offline versions of this method for a long time. A good example of an offline A/B-like test is the bean growing experiment commonly practiced in school – a jar is placed in a dark place, while another is placed somewhere with a fair amount of Sun exposure, and the growth results are compared. In this case, the variable whose effects are being studied is "Sunlight" (Armour, 2015). Returning to online, these alternative versions could range from minimal aesthetical alterations, such as changing the colour of a button, to a complete layout revamp.

As illustrated in Figure 3, the test itself works by splitting the page's visitors through the several alternatives being tested (e.g. one set of visitors sees alternative A, the other B, …), and then calculating and comparing the relevant performance metrics (e.g. click-through rates, conversion rates, etc.) across them.

Figure 3 - Representation of an A/B test

Kohavi and Thomke (2017b) refer a particular case at Bing as an example of the benefits that such tests can present. It was 2012 and an employee had an idea regarding a change in the way the search engine presented their advertisements. Even though that idea required minimal effort, it was one among hundreds, which led to its relocation to the bottom of the priority list. After 6 months, nevertheless, when an engineer saw that the cost for that change was nearly inexistent, he launched an A/B test to assess its impact. Surprisingly, that small change generated a significant increase of 12% in revenue, which boiled down to $100 million a year in the United States alone. Given its small proportion, no one would have thought that such a change would have an impact that noticeable: without experimentation, Bing would have probably wasted a great opportunity. This demonstrates the importance of having the capability to run several tests – the higher the amount performed, the higher the probability of hitting a jackpot (Kohavi & Thomke, 2017c).

Nowadays, the most successful tech giants – including Microsoft, Google and Amazon – run over 10000 online controlled experiments on an annual basis. Figure 4 represents the increasing use of experimentation by Bing over the years.

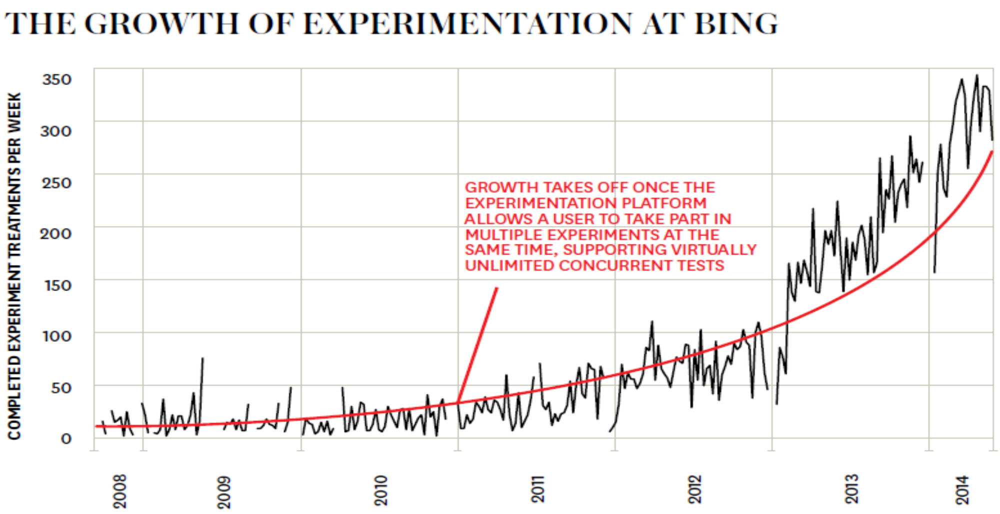

Figure 4 - Bing's growth in experimentation

3 Problem Description

The main objective of this study is to analyse the impact of the improvement of certain fashion-related imagery variables on the website's performance metrics, such as the Conversion Rate (CR).

Before proceeding to the analysis, it is relevant to understand what exactly is an imagery variable in the fashion industry.

3.1 Variables that influence Fashion Imagery

As previously mentioned, the first step towards enhancing the quality of fashion imagery is understanding the variables that influence it.

When someone visits a fashion e-commerce website, such as Farfetch.com, their main sources of imagery exposure are product listing pages (PLPs), product display pages (PDPs), and finally, editorials. The aspects of imagery that concern PLPs are different than those that concern PDPs. Regarding editorials, the relevant variables are a mix of those relevant to PLPs and PDPs. Having established this, it makes sense to divide the variables into PLP and PDP-related groups.

On the PLP level, where several products are showcased together, the first variable to consider is the image's scale. As several products are displayed side by side, it is very important to make their images relate to the same scale in order to provide the customer with a comparison point. Also important is product alignment, as it not only facilitates the aforementioned comparation of products but also provides a more aesthetically pleasing experience. Not so related with the images' content itself, it is important to adjust the number of images displayed by row and column, as well as their size, to the website's general appearance. Last but not least, hover content (whatever extra content is displayed when a customer hovers its mouse on an image) is also something to consider. Since not all customers proceed to a product's display page straight away, hover content presents itself as another chance of triggering the visitors' curiosity to check the products' pages.

On the PDP level, on the other hand, the spotlight is over only one product at a time, which means that comparison-affecting variables are not relevant. In this case, the most important variables are the ones that promote the product and its quality, and these are subdivided into three groups: physical variables, digital variables and external assets. The first physical components are the existence (or not) and exposure level of a model. Featuring products on a model has a great positive impact on sales (Cardew, n.d.), as a bond is created with customers – by seeing someone wear an item they more easily imagine themselves wearing it too. "Exposure level" is self-explanatory. It is different to feature a full body model than only his/her chest, feet, etc. The presence of a model, in turn, results in new variables being added to the content. First, there is styling. It is of paramount importance to create the best combinations of clothes possible. This will not only potentiate the sales of the featured product, but possibly even lead customers to purchase the adjacent products used in the combinations. In Farfetch's particular case, styling is very dependent of the toolkit (the set of items used by stylists to create outfits). The greater the depth of the toolkit, the more possible combinations can be made, which means the probability of achieving great combinations is much higher. Still relating to the model variable, the poses adopted when on-set will communicate a lot about the products' nature (Pegram, 2008).

Apart from the model and its aforementioned related variables, the choice of background will also have a great impact on customer perception. While the standard white background is known to be very adequate for fashion photography, adopting patterns, colours or even pictures could drive sales, depending on the product, by further inspiring the customer. Next, lighting is incredibly important. If not correctly adjusted, wrong colours could be shown to the customer, as well as small wrinkles. If carefully chosen, on the other hand, it could help show the product in detail, and potentially "give it more life".

With regards to the digital components, one critical item is the resolution of the images. There is a slight trade-off between quality and performance - the higher the resolution (quality), the heavier the image - which negatively affects its loading time. Resolution also impacts the degree to which a customer can zoom-in on a product (to see its details). If a lower resolution has to be used, extra product shots of the product could be taken in order to display the necessary details to the customers. This is not the only advantage of utilizing extra shots, as they can also be used to highlight specific components of the product, such as its selling points.

The use of 360 images and/or videos is also much discussed. On one hand, these assets can be very useful to exhibit the products' details and engage the customer, which expectedly rises the intention to buy. On the other hand, however, these components represent a higher production cost, as well as extra process complexity.

Finally, regarding external assets, runway imagery and user generated content can be considered. Runway imagery is expected to have a positive impact on the customer's inspiration to buy, as the products are showcased by famous models on the world's most renowned runways, such as the Paris Fashion Week. User generated content also has a positive impact on inspiration, as well as strengthening the bond between the seller and the buyer – seeing other regular people wearing the products purchased on the website creates a sense of trust.

3.2 Farfetch's production

Since this study involves some alterations in the production guidelines of certain product types, it is important to include an overview of Farfetch's "as is" production process.

Boutiques and brands from all over the world add new products online and send them to one of Farfetch's production centers (according to their location) to be photographed for the website. The production process starts with the arrival of these products, in slots, at the production center.

The items follow a standardized flow, starting at the Scan-In section, where their arrival is registered in the production database. It is also here that duplicates (items that were already produced) are detected and blocked. All items that progress through Scan-In are organized and placed in rails. These rails then travel through the rest of the process.

Currently, there are 3 main types of photography in Farfetch: Flats, Stills and Live Model. All item families feature either flats or stills photos. Live Model, on the other hand, is only applied to Clothing and Bags.

Depending on the item family, a given product can have different numbers of photos. For instance, Clothing items usually have 1 Flat and 4 Live Model shots, Bags have 1 Live Model and 4 Still shots, Shoes have 4 Still shots, Accessories have 3 Still shots and, finally, Jewellery features 2 Stills and 1 Live Model shot.

Items to be photographed on live model are sent to the ironing station, where wrinkles and eventual small stains are removed. After the stylists finish combining outfits for the items (making use of other items on the rail and the toolkit), the items are photographed on a model (each item takes around 3:30 min to be completed). Next, the items are forwarded to the Stills/Flats studios, where they take the remaining needed photos according to their category. Items that are not featured on a model, on the other hand, are sent straight to Stills/Flats.

From this point on, there are two different flows to track: the flow of the item throughout the physical stations of the production centre and the flow of the images through the digital tracks of the production process.

In the case of the digital tracks, right after photos are taken, they go through a very strict quality control (QC) that guarantees that the lighting and the outfit as well as the model's pose are correct in the case of photos taken on live model. If the images fail the QC test, there is a high likelihood that the item needs to be shot again.

After proceeding through the QC checkpoint, the photos are edited by the Post Production team. It is at this station that all studio props (such as mannequins) are removed from the images, and where they are given the final touches before being uploaded to the server. It is important to have in mind that this upload does not mean that the item is automatically online (available on the website). In order to go online, several other aspects need to be guaranteed, such as the existence of a description for the product and a stock value greater than 0 in any store.

While the images follow the aforementioned track, the physical products are scanned-out of the production centre, packed and shipped back to the partners.

3.2.1 Production performance metrics

Similarly to any other process, it is important to constantly assess Farfetch's production performance.

Although the final product is in a digital form, Farfetch's production behaves very much like any factory, and so the metrics used to assess its performance are also very similar. For instance, the "daily production" consists of the number of products photographed in a day's time. In terms of lead time, two separate metrics are considered: "time to return slot" measures the time it takes for all the products inside a slot to travel through the whole production process (scan-in to scan-out) and "production lead time" determines the time it takes from the moment a product is scanned-in to its upload. "Time to return slot" is especially important for the partners, as it directly relates to the time they have to wait to get their product back, which ultimately represents the time their stock unit is unavailable.

4 Methodology

This chapter provides an outline of the methodology used to assess the subject of research.

It includes the variable and target product category choice decision processes, the data collection method, the performance metrics taken into account and the statistical validation methods used to validate the results obtained.

4.1 Variable choice

Ideally, results would be presented for each, and every one of the aforementioned imagery variables. However, due to the thoroughness needed to achieve accurate insights and the limited time for the study, only one variable is tested.

Since there are so many imagery variables to consider, picking one to place under the microscope can be somewhat challenging.

One of Farfetch's main objectives for the upcoming years is reducing its return rate. While visual content is mainly perceived to impact conversion, some image variables can also, theoretically, have an impact on returns. This is the case of variables that influence the level of product detail presented – the better the customer knows what he/she is about to buy, the lower the probability of an unpleasant surprise upon the reception of the product.

Given the returns reduction goal, it seems logical to pick one of these variables and, therefore, we choose the extra shots. By taking close-up photos of a product, revealing the fabric's complexion, hidden pockets, fancy stitching, etc. the probability of a customer being surprised upon the arrival of its order (either positively or negatively) is reduced, which means that the probability of return is, theoretically, lower. Furthermore, some specific selling points of the product can be better highlighted, which expectedly enhances the inspiration and intention to purchase.

4.2 The A/B Test

Through testing we are able to obtain accurate insights on the value associated with a change, without incurring in unnecessary costs. This, in turn, allows us to make informed decisions of whether or not to proceed to the full rollout of the project.

In order to determine the actual value of featuring extra shots, an A/B test is set for 6 months. During this time, each product selected for the test is uploaded to the website with two different versions: control, no extra shots; and alternative, with 5 additional shots. Each time someone lands on one of these products' display pages that person will automatically be shown either the control (regular number of shots – Figures 5 and 7) or the alternative version (regular shots plus 5 extra views – Figures 6 and 8). Henceforth, control and alternative version will be called "A version" and "B version", respectively. In order to guarantee visual consistency, when a new visitor is presented with either the A or B version, all the product pages he/she visits afterwards will reflect that same version.

Figure 5 - The five original item photos

Figure 6 - The five original photos, plus five extra shots

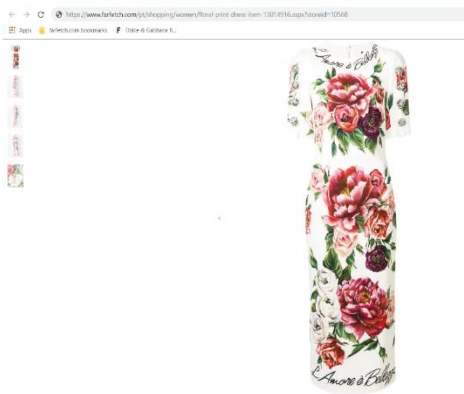

Figure 7 - Original product photo display page

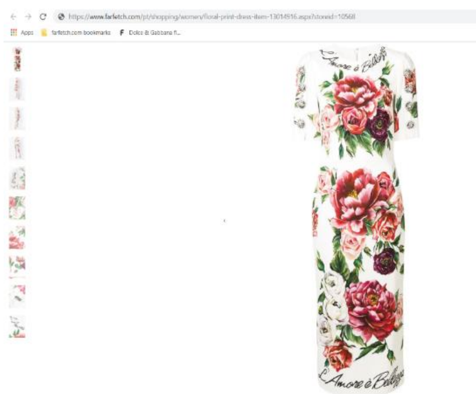

Figure 8 - Product photo display page with extra shots

4.2.1 Target Product Categories

In order to achieve statistical significance on the A/B test, we not only need to have a wide enough range of products, but also enough number of visits on both A and B sides. As indicated in Figure 9, Clothing accounts for around 57% of the products that arrive in Guimarães every year, which means that the desired number of products for the test would be quickly reached.

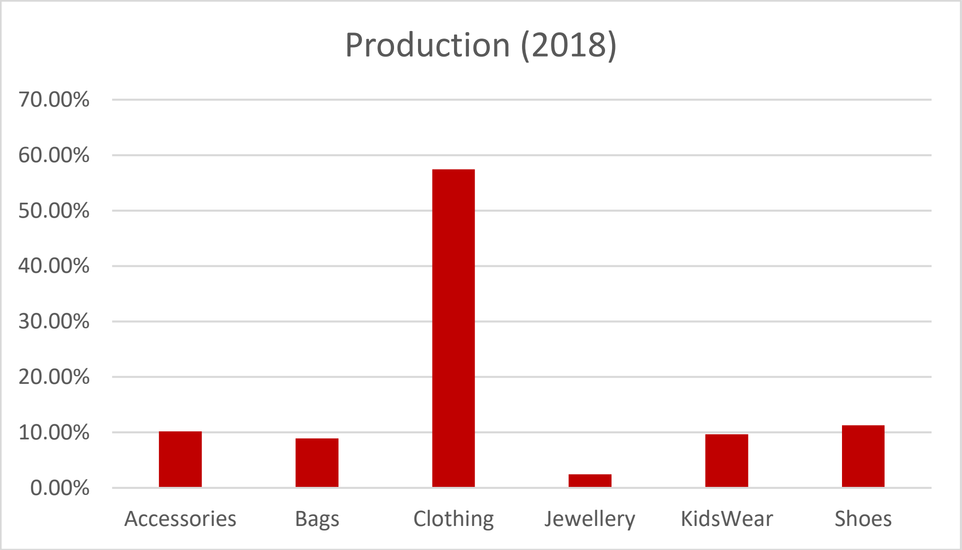

Figure 9 - Farfetch's 2018 production percentage by product category

Furthermore, and most likely due to the higher volume of products, this category also exhibits the highest share of visits (which can be seen in Figure 10), which is the second most important success factor for the A/B test. Taking into account the production capacity, the threshold for the minimum number of products was set at 6 thousand.

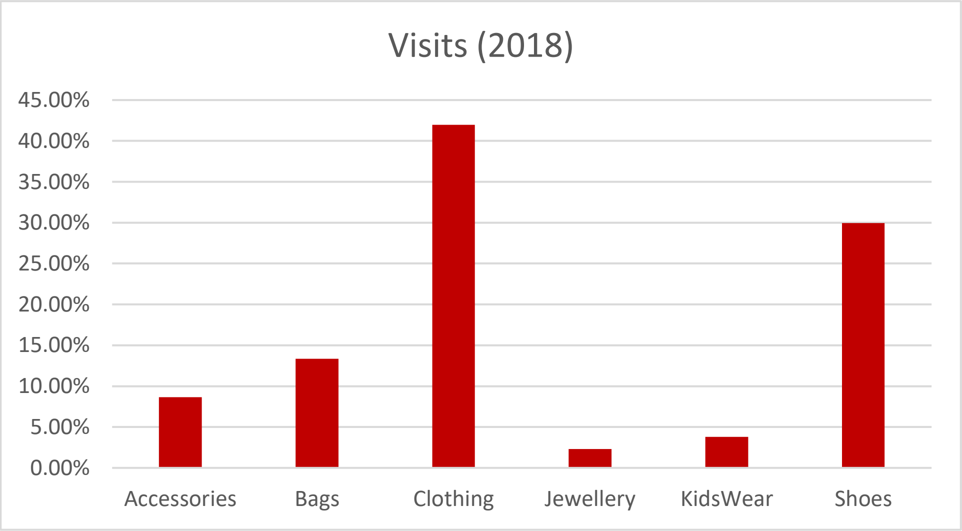

Figure 10 - Farfetch's 2018 traffic volume by product category (percentage)

Since the improvement being analysed is the addition of product shots, the number of shots to add must also be defined. After discussing the topic with Creative Services (Creative Operations' team responsible for the overall creative direction and look & feel of the website), 5 extra shots is, according to the team, sufficient to showcase all the important details of the products.

The last pre-analysis decision concerns the number of products. Taking into account the predicted increment in unitary production time and the spare production capacity, the product volume threshold is set at 6000.

Although the planned product volume was 6000, due to some technical problems and the naturally occurring error in process impact calculations, only 5800 items were actually produced.

4.2.2 Measuring variations in performance metrics

Before commencing the analysis, the dataset needs to be streamlined. In this case, this means removing all data regarding items that had no visits on either their A or B side. By applying this filter, 800 products are eliminated, reducing the sample to 5000.

This analysis consists on the calculation of the uplift verified on the comparison between the A and B versions in each of the following metrics: Add to Wishlist Rate (AW2R), Add to Bag Rate (A2BR), Started Checkout Rate (SCR), Conversion Rate (CR) and Return Rate (RR). The values measured by each metric are described in equations (4.1) to (4.5).

$$A2WR = \frac{Sessions\ that\ Added\ a\ Product\ to\ the\ Wishlist}{Total\ Number\ of\ Sessions} \tag{4.1}$$

$$A2BR = \frac{Sessions\ that\ Added\ a\ Product\ to\ the\ Bag}{Total\ Number\ of\ Sessions} \tag{4.2}$$

$$SCR = \frac{Sessions\ that\ Initiated\ the\ Checkout\ Process}{Total\ Number\ of\ Sessions} \tag{4.3}$$

$$CR = \frac{Number\ of\ Orders}{Total\ Number\ of\ Sessions} \tag{4.4}$$

$$RR = \frac{Number\ of\ Orders\ Returned}{Total\ Number\ of\ Orders} \tag{4.5}$$

Furthermore, all these metrics, except the A2WR, relate to each other according to the customer funnel represented in Figure 11. Please note that the proportions are not accurate.

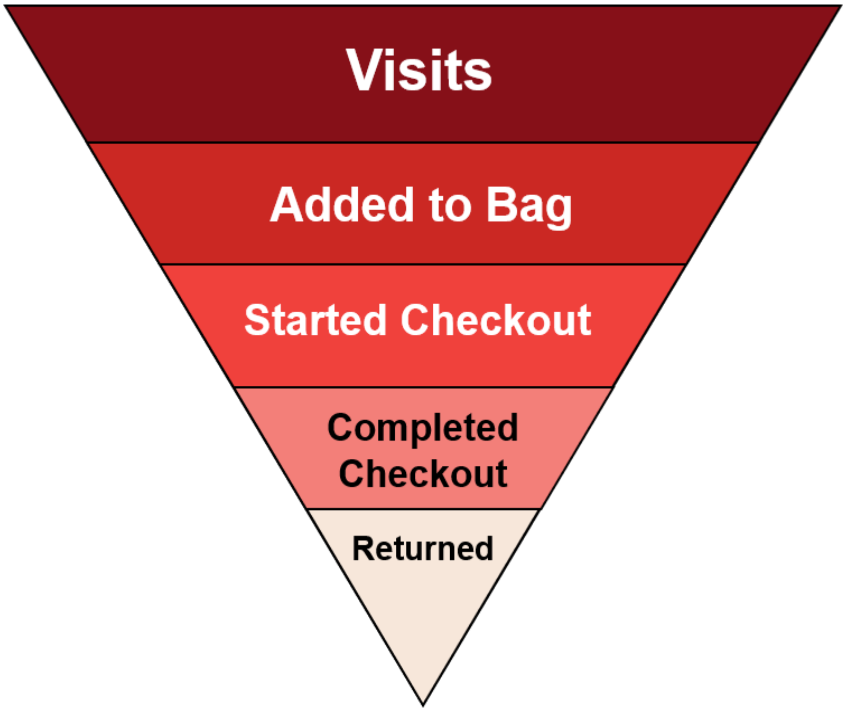

Figure 11 - Customer journey funnel representation

As “Added to Wishlist” doesn’t have a direct relation with any of the other figures, it is excluded from the funnel. Analysing so many metrics that relate directly to each other might seem slightly redundant but is, in fact, very useful, as it is easier to achieve significant results with metrics that have higher values (such as A2BR in comparison to SCR and CR). For example, assuming no other conditions throughout the customers’ journey vary, a given uplift in A2BR will be exactly the same as in CR - if there is a percentual increase in the number of “add to bags”, a corresponding percentual increase in conversions is expected.

Assuming X represents any of the aforementioned metrics, the uplift calculation is made through the use of (4.6):

$$Uplift_{B,A} = \frac{X_B - X_A}{X_A} \quad (4.6)$$

It is important to keep in mind that positive uplifts in A2WR, A2BR, SCR and CR are considered favourable results, while the opposite is true for RR.

Not all uplifts can be taken into account, however. In order to refine the obtained insights on such an analysis, a significance test must also be performed. Since the uplift can be either negative or positive, the two-tailed, two-proportioned z-test is used:

$$H_0 : X_B = X_A$$
$$H_1 : X_B \neq X_A$$

First, the difference between B and A values for the metric is calculated as in (4.7):

$$Diff_{B,A} = X_B - X_A \tag{4.7}$$

Since we are testing the validity of the null hypothesis, we must assume that $X_A = X_B$, when calculating the standard error and, therefore, define X as the weighted average of X_A and X_B, as indicated in (4.8).

$$X = \frac{Visits_A \times X_A + Visits_B \times X_B}{Visits_A + Visits_B} \tag{4.8}$$

Consequently, the standard error is determined by (4.9):

$$SE_{X_B - X_A} = \sqrt{\left(\frac{X_A \times (1 - X_A)}{Visits_A} + \frac{X_B \times (1 - X_B)}{Visits_B}\right)} = \sqrt{\left(\frac{X \times (1 - X)}{Visits_A} + \frac{X \times (1 - X)}{Visits_B}\right)} \tag{4.9}$$

As demonstrated in (4.10), by dividing the Difference by the Standard Error, the z-score is obtained:

$$z - score = \frac{Diff_{B,A}}{SE_{X_B - X_A}} \tag{4.10}$$

Finally, the p value for the z-score is determined. As the test is being done at a confidence level of at least 95%, only the differences with associated p values of 0.05 or less can be considered statistically significant. P value lower than 0.01 indicate a confidence of at least 99%.

The lower and upper bounds for the confidence interval of the difference are calculated as indicated in (4.11) and (4.12) with a confidence level of 95%. All non-significant differences are contained in an interval that includes 0 (as the one in Figure 12) which further corroborates their non-significance – we can't be sure, at a 95% confidence level, that a given uplift is in fact positive or negative.

$$LB = Diff_{B,A} - 1.96 \times SE_{Diff_{B,A}} \tag{4.11}$$

$$UB = Diff_{B,A} + 1.96 \times SE_{Diff_{B,A}} \tag{4.12}$$

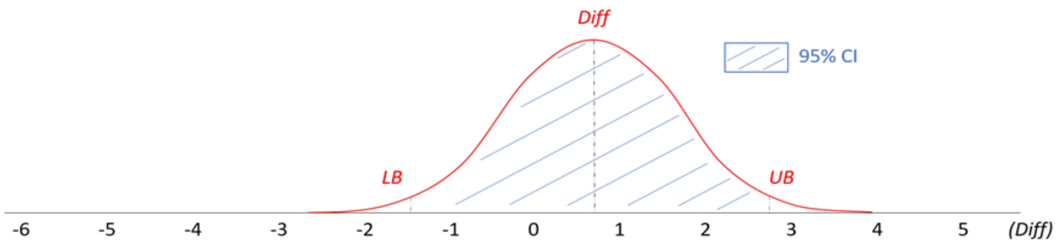

Figure 12 - Representation of a measured difference's range of values

An initial analysis was done on the impact of the extra shots from a global point of view.

This global analysis was followed by five detailed analysis on customer type, device, country. product price-point and product sub-category.

4.3 Measuring the impacts on process efficiency

Since taking ten photographs of a product is more time-consuming than taking only five, before proceeding to the rollout decision, one must understand the effects on the efficiency of the production process.

The production cost per item is determined by dividing the daily wages paid to every employee involved in the production process by the number of items produced per day (a slight increase in this metric is expected). This variation is determined as follows:

$$\Delta Cost\ Per\ Item = \frac{\sum Daily\ Wages}{Daily\ Production_F} - \frac{\sum Daily\ Wages}{Daily\ Production_S} \tag{4.13}$$

Where:

F stands for “Finish” and represents the values after implementation of extra shots and

S stands for “Start” and represents the values before implementation of extra shots

4.4 Rollout decision

After gathering all the aforementioned information, a decision has to be made regarding the full implementation, or not, of extra shots.

Two rollout scenarios are defined, for which the added costs and ROI are calculated. The increase in costs is calculated by extrapolating the increase in the production cost per item to the expected production volumes of both scenarios:

$$Added\ Costs = \Delta Cost\ Per\ Item \times Production\ Volume \tag{4.14}$$

Regarding ROI, the calculations are slightly more complex. First, the number of visits, orders and the average order value of the product types considered in each scenario are calculated. The number of visits is then waged against the global number of visits to Farfetch in order to estimate the reach of the products.

$$Reach = \frac{Number\ of\ Visits_{Scn}}{Total\ Number\ of\ Visits} \tag{4.15}$$

The reach represents the % of business volume covered by the scenario, and is used to extrapolate the uplift on the scenario's CR to Farfetch's global CR, through the following formula:

$$\Delta CR_{FF} = \frac{\left[\frac{CR'_{Scn} \times Reach'_{Scn} + [CR_{FF} - (CR_{Scn} \times Reach_{Scn})]}{(1 - Reach_{Scn}) \times (1 - CR_{Scn})}\right]}{\frac{CR_{FF}}{1}} - 1 \tag{4.16}$$

Where:

$\boldsymbol{CR_{FF}}$ is Farfetch's global conversion rate,

$\boldsymbol{CR_{Scn}}$ is the conversion rate of the product group considered in a given scenario and

$\boldsymbol{CR_{Scn}}'$ is the uplifted conversion rate of the product group considered in a given scenario

Since the uplift measured measured in each scenario's conversion rate has a 95% confidence interval, the extrapolated uplift in Farfetch's global conversion consists of an interval of possible values rather than a single one. This uplift is then multiplied by the forecast of the baseline sales volume for the following year (GMV) in order to determine the ROI:

$$ROI = \Delta CR \times GMV_{BaselineForecast} \tag{4.17}$$

Again, since the extrapolated uplift is a range of values, so will the ROI be too.

Before proceeding to the roll out scenario decision, a final analysis is done to determine the effective number of extra shots needed. This consists in determining the percentage of visitors that see each number of photos (1, 2, 3, …) by device (Desktop, Mobile and Tablet).

Based on these calculations and production capacity, one of the scenarios is chosen.

4.5 Monitoring Framework

Although the test indicated very positive results for the implementation of extra shots, we need to have in mind that the products, customers and market tendencies keep shifting every year, which means that the effective uplift guaranteed with the implementation is not necessarily the same as the one predicted through testing. That being said, it is important to keep track of the effective impact of the improvement, and for that purpose a monitoring framework was built.

This monitoring framework consists of a set of queries that after implementation will be periodically ran on Farfetch’s databases in order to assess the performance of the item types involved (for products belonging to sub-categories “Coats”, “Denim”, “Dresses” and “Trousers”). This framework’s working principle is the comparison of products affected by the aforementioned changes with products of the same type that were not altered. The selection of comparable products to analyse presents itself as the biggest challenge in this monitoring framework, as sometimes the definition of “comparable” is somewhat not clear.

In this case, we decided to take a hybrid approach, using similar timeframes and product volumes for comparison. For that effect, we compare the conversion and return rates of the items included in the aforementioned sub-categories (produced after the definitive rollout date of the enhancement) with those of the equivalent types of product produced (before the enhancement took place). Furthermore, in order to avoid biased readings, two additional measures must be taken. First, the performance of the sets of items is assessed during equivalent time periods – enhanced items are evaluated on a timeframe starting on the extra shots rollout date and non-enhanced items are assessed on the same timeframe in 2018. Secondly, only full-price sales are taken into account, as discounts usually have a great impact in both metrics.

On an initial phase, these comparisons will be made on a biweekly basis, allowing a close and strict monitoring of the improvement. Once it is determined that the results are minimally stable and consistent, the periodicity can be gradually reduced – first to a monthly basis, followed by bimonthly and, finally, annual.

In the unlikely case that the results are not positive, production of the selected type of items can return to the previous guidelines.

5 Results

This chapter is dedicated to the presentation and discussion of the key results obtained in the investigations mentioned in the previous chapter, conducted during the A/B test. Complete tables with all the results can be found in the appendix.

5.1 Global Analysis results

The first analysis is done on a global level – this means an overall performance comparison of the alternative versus the control. Potentially the most valuable analysis, due to the fact that almost no sample filters were used – making the sample larger and more representative of the population – there is a high expectation as to what insights can be retrieved from it.

Much to our dismay, no significant results are observed in this analysis. The measured -0.7420%, 0.2772%, -0.1189%, 1.3274% and 1.2483% uplifts in A2WR, A2BR, SCR, CR and RR – respectively – are cast aside as non-significant with the corresponding p values (95% confidence level) of 0.4226, 0.6909, 0.9259, 0.4995 and 0.7500 (Figure 13).

Global - Add to Wishlist Rate					
	Uplift	**Diff**	**SE**	**z-score**	**p value**
Alternative	-0.7420%	-0.0389%	0.0485%	-0.8016	0.4226
Global - Add to Bag Rate					
	Uplift	**Diff**	**SE**	**z-score**	**p value**
Alternative	0.2772%	0.0247%	0.0622%	0.3976	0.6909
Global - Started Checkout Rate					
	Uplift	**Diff**	**SE**	**z-score**	**p value**
Alternative	-0.1189%	-0.0034%	0.0362%	-0.0931	0.9259
Global - Conversion Rate					
	Uplift	**Diff**	**SE**	**z-score**	**p value**
Alternative	1.3274%	0.0163%	0.0241%	0.6752	0.4995
Global - Return Rate					
	Uplift	**Diff**	**SE**	**z-score**	**p value**
Alternative	1.2483%	0.2551%	0.7751%	0.3291	0.75

Figure 13 - Global Analysis Results

Note: *, ** and *** represent statistical significance at a 95%, 99% and 99.9% level, respectively

5.2 Deep-dive analysis results

Proceeding to a thorough analysis on several factors, which involve the division of customers into groups according to type (returning or new), channel/device, location, product selling price and product sub-category (detailed product classification within the Clothing category).

Starting with the customer type, the only significant uplift measured is -21.1% on the RR of new customers (Figure 14), value associated with a p value of 0.0294. Since the return rate of new users is usually higher than that of returning ones it is, in theory, easier to provoke its variation – hypothesis corroborated by the results obtained.

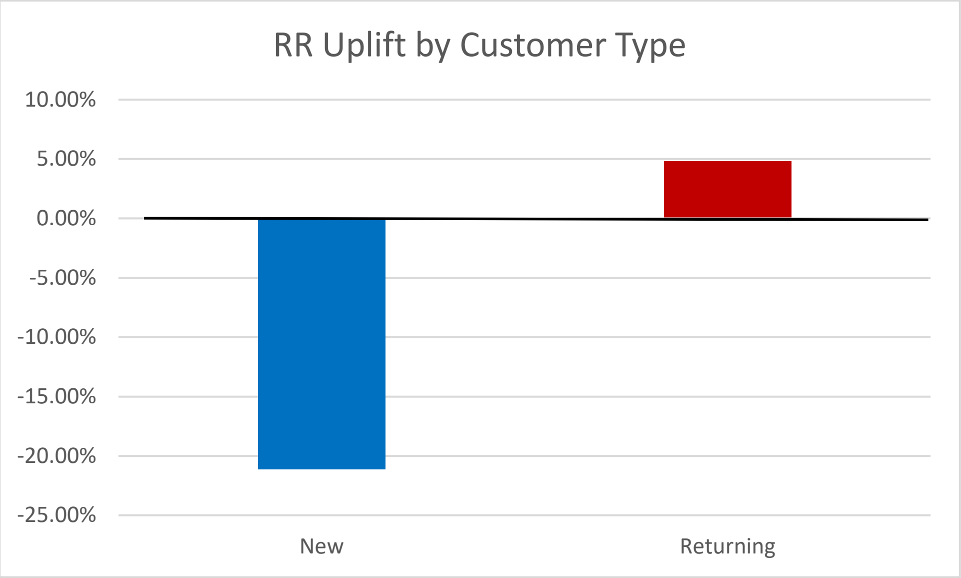

Figure 14 - Return Rate uplift by type of customer

The second study is done for each device (desktop, mobile and tablet). A clear trend is observed here, with high uplifts in all metrics (except A2WR) for tablet users. With a p value of 0.0027, the latter has a significant uplift of 18.0% in SCR. The significance of this number is validated in Figure 15, where it is visible that among all device types, only Tablet has both the lower and upper bounds of its SCR difference confidence interval above 0. This is possibly explained by the fact that the user interface for Tablet users is much more visually oriented and encourages interaction with the photos. Actually, at the end of the chapter, the chart present in Figure 18 provides proof that customers on this channel interact more with the images.

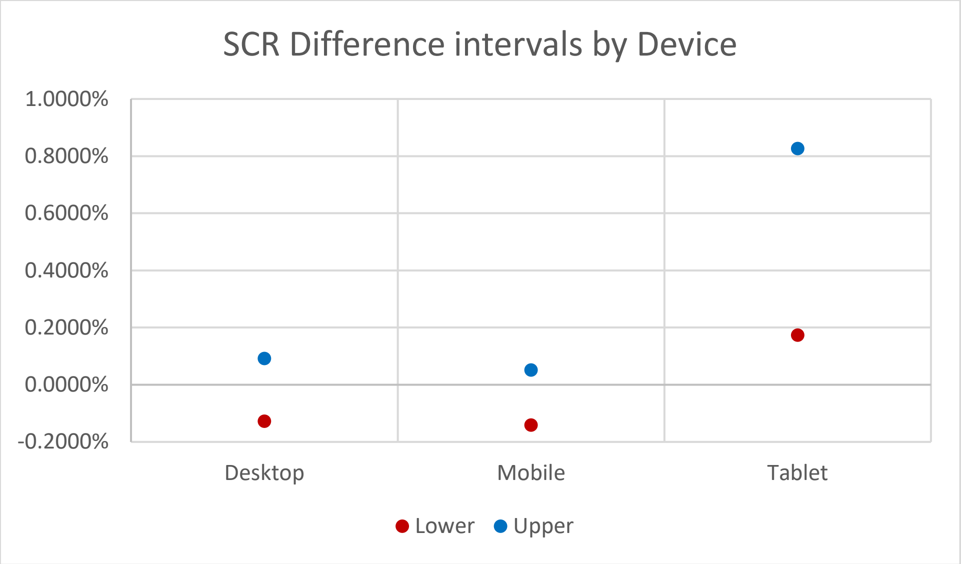

Figure 15 - Range of values for the measured differences in the "Started Checkout Rate" metric

The third analysis is done at a country level, and the majority of the uplifts calculated are not significant. It is important, however, to highlight the significant uplifts of 71.4% and -37.7% in RR in China and France (respectively) which are visible in Figure 16. It is important to reinforce the fact that positive uplifts in Return Rate are undesired.

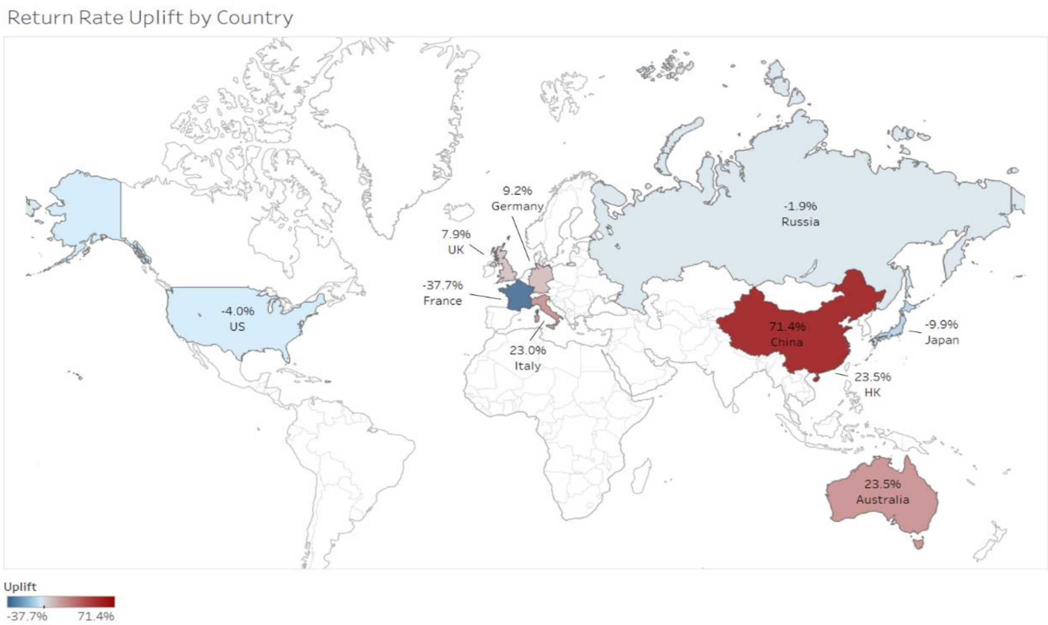

Figure 16 - Return Rate uplift by country

In relation to the price-point, no considerable effects were detected.

Finally, pertaining to the product sub-categories analysis, four significant uplifts of 12.2143%, 9.5297%, 6.5412% and 6.1266% are verified for Coats, Denim, Dresses and

Trousers, respectively. The significance of these figures is corroborated by their respective p values of 0.0037, 0.0072, 0.0002 and 0.0197, as seen in Figure 17.

Add to Bag Rate					
Sub-Category					
Coats	**Uplift**	**Diff**	**SE**	**z-score**	**p value**
Alternative	12.2143%	0.5511%	0.1900%	2.9005	0.0037**
Denim	**Uplift**	**Diff**	**SE**	**z-score**	**p value**
Alternative	9.5297%	0.4714%	0.1755%	2.6860	0.0072**
Dresses	**Uplift**	**Diff**	**SE**	**z-score**	**p value**
Alternative	6.5412%	0.3529%	0.0946%	3.7304	0.0002***
Trousers	**Uplift**	**Diff**	**SE**	**z-score**	**p value**
Alternative	6.1266%	0.3018%	0.1294%	2.3323	0.0197*

Figure 17 - Add to Bag Rate uplift by product sub-category

Note: *, ** and *** represent statistical significance at a 95%, 99% and 99.9% level, respectively

5.3 Decision to proceed to rollout

After completing the impact analysis of the A/B test, it is time to decide whether or not to proceed to the full implementation and, if so, what is the most adequate way of executing.

5.3.1 Scenario definition

In order to avoid the overlap of results (e.g. the uplifts calculated on the referred sub-categories including part of the uplift verified for tablet users), only the results obtained in a single scope are used to build the business case.

Since the most significant, consistent and actionable results were obtained in the sub-category analysis, two scenarios are defined for the rollout.

Scenario 1 is the production of all coats, denim, dresses and trousers, and the second scenario (Scenario 2) is the production of the same subcategories but only for the top 50 brands in terms of sales volume.

5.3.2 ROI and profitability

As mentioned in Section 4, the first step towards understanding the potential returns is the extrapolation of each scenario-specific uplift in CR (ΔCR_{Scn}) to Farfetch's global conversion rate (ΔCR_{FF}). Through these calculations we can infer that Scenario 1 could represent an uplift of 0.24% to 1.88% in Farfetch's global annual conversion rate, which directly translates into an extra 0.24% to 1.87% in GMV.

Scenario 2, in turn, could mean an additional 0.11% to 0.81% in CR or GMV.

5.3.3 Number of shots

After assessing the costs, benefits and feasibility of each scenario, it is also very important to verify if there is a need to iterate the number of extra shots. For that purpose, we evaluated the number of photos seen by product across all devices. This is accomplished by counting the number of right clicks/swipes on photos. In the particular case of Desktop, visitors are shown a pair of photos at a time, instead of a single view. Each time they right click (to scroll through the photos), only one of the photos of the pair varies. Therefore, two right-clicks on a Desktop PDP is equivalent to three right-swipes on a Mobile/Table PDP. Considering this, the chart in Figure 18 is obtained:

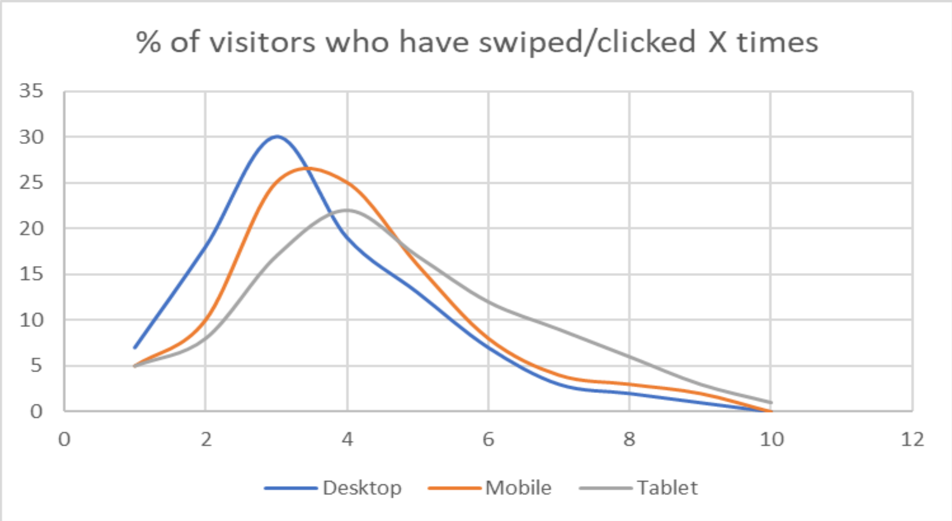

Figure 18 - Percentage of visitors seeing each number of photos

Figure 18 depicts the percentage of visitors that clicked/swiped each amount of times by device. Apparently, tablet users swipe more than both desktop and mobile users, seeing approximately seven photos. Having in mind that the default number of photos is five, this proves that there is no need to produce more than two extra shots.

5.3.4 Feasibility and final decision

In theory, by sacrificing more or less process speed in some of the studios, it is possible to adopt any scenario. Nevertheless, since it is crucial to ensure that the products are returned to the partners in a timely fashion (time to return slot metric consistency), the US and Hong Kong studios would need to be expanded (in order to maintain the lead times under control) if Scenario 1 was adopted.

Having in mind the extra costs associated with the expansion and the added complexity, the final decision is to proceed with the rollout of Scenario 2.

6 Next Steps and Conclusions

The conducted experiment ended with very positive results and the decision to implement extra shots in the aforementioned product sub-categories across the top 50 brands, which is expected to have an excellent impact in revenue. The immediate next steps consist of an extension of the analysis done, as well as well as ensuring the successful implementation of the extra shots on the selected type of products.

6.1 Next steps

Regarding the study itself, some questions were left unanswered, such as the reasoning behind the substantial uplift in China's Return Rate upon exposure to extra product shots. Market studies must be done in order to explain some of the results obtained. Furthermore, an extension of the study could be made to test the impact of extra shots on other product categories, such as Shoes or Bags.

Pertaining to the definitive implementation of extra product shots according to Scenario 2's criteria, certain conditions must be guaranteed in order to achieve success. First, a digital tool that highlights eligible products when scanned at the Scan-In station must be developed to enable the technicians to identify the product with a special tag that notifies the photographers of the need to shoot the extra views its extra views. Secondly, since only 2 extra shots will be produced – as opposed to the 5 used during the test – certain photography guidelines must be re-defined. Finally, adequate adjustments must be done to the photographers daily production KPIs, taking into account the extra time and effort needed to shoot the 2 extra photographs.

It is also important to keep in mind the possibility of adopting Scenario 1 (which represents a greater reward than Scenario 2) on the long run, which means that all requirements to guarantee that possibility must be thoroughly determined.

6.2 Critical reflection

Overall, the experiment was considered very positive. The test, although relatively cheap, led to valuable insights on a great business opportunity, which further reinforces the importance and value of testing in an industry like Farfetch's. The methodology used can, with the adequate adjustments, be replicated in further similar studies, which means that the company's testing flexibility was slightly increased. Finally, as the obtained results were so positive, awareness was raised inside Farfetch towards the importance of imagery and the opportunity it's excellence could represent.

Even more important than praising the positive results, is to adopt a critical stance regarding what could have been done in a different, perhaps more correct, way. The most noticeable negative aspect was the fact that the added shots were not used, to a certain extent, to their full potential. As proved by the analysis of the number of photos seen by the customers, typically only 7 out of the 10 total available photos were seen. As seen in Figure 19, when a customer visits a product page he/she is not made aware, right away, of the number of distinct views of the item.

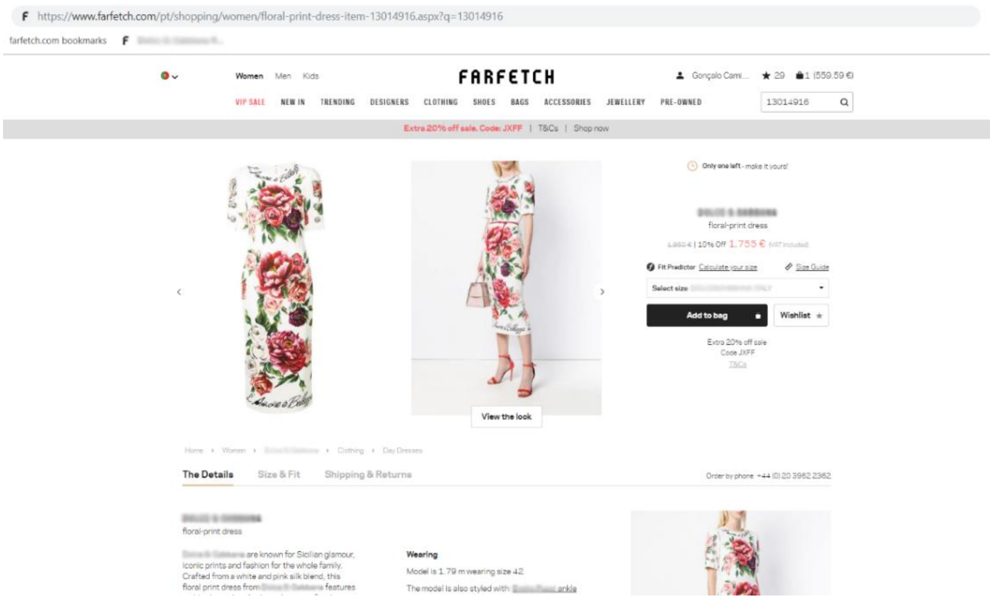

Figure 19 - Product display page (as is)

Perhaps, if visitors knew, beforehand, all the images they could explore – through thumbnails, for example (Figure 20) -, a higher level of engagement could have been achieved which could, in turn, have led to even better results.

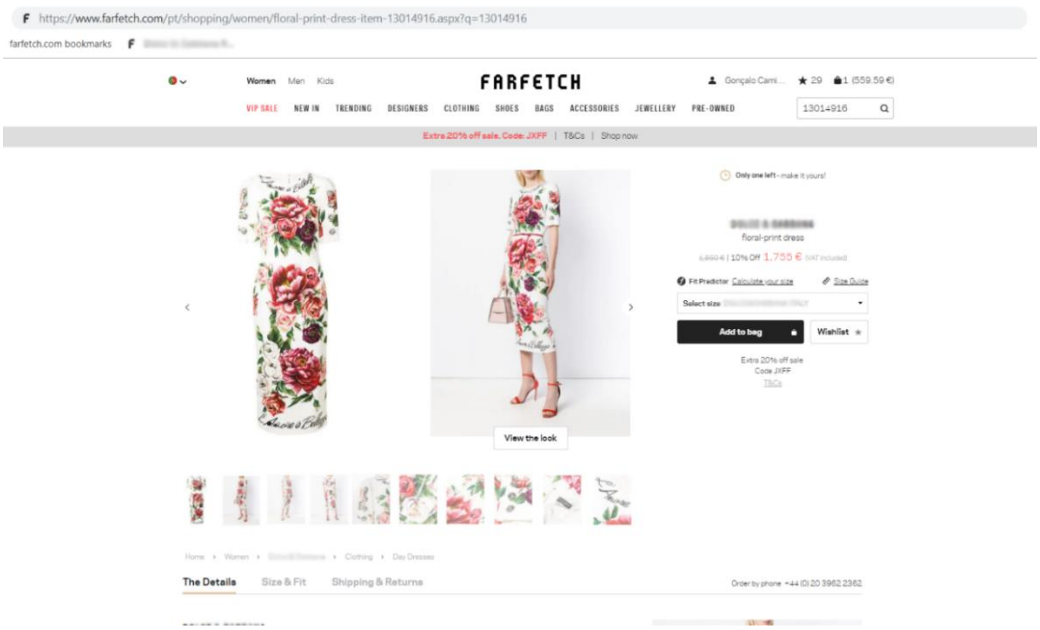

Figure 20 - Concept of a product display page with photo thumbnails

Perhaps, when doing further analysis of this kind, we could spend some extra time determining how can the assets being tested be used to their maximum potential. Finally, additional analysis that were not done during the A/B test, nor are included in the monitoring framework - such as the impact of the extra shots on customer retention – could be done.

6.3 Further experiments

As seen in Chapter 3, Farfetch still has a lot of variables in need of an enhancement impact test such as this one.

First, a criteria must be defined in order to prioritize which variables shall be tested first.

After having a very clear idea of what must be investigated first, the methodology used in this test could be adapted to test the other variables.

Reference List

Adii Pienaar 2015, What is A/B Testing, accessed 15 June 2019, < https://conversio.com/blog/ab-testing-for-ecommerce/>

Armour, H. (2015). A / B Testing. *Digital Marketing*, (September 2015), 201–208. https://doi.org/10.1016/B978-0-12-404619-1.00013-7

Bain & Company. (2019). The Future of Luxury: A Look into Tomorrow to Understand Today. *Bain & Company Milan*. Retrieved from https://www.bain.com/insights/luxury-goods-worldwide-market-study-fall-winter-2018/

BAIN & COMPANY Media Center. (2018). *The personal luxury goods market delivers positive growth in 2018 to reach 260 billion-a trend that is expected to continue through 2025*. 7.

BCG, & Altagamma. (2018). *The Key Trends for the Luxury Goods Industry in 2018*.

Bui, M., & Kemp, E. (2013). E-tail emotion regulation: Examining online hedonic product purchases. *International Journal of Retail and Distribution Management*, *41*(2), 155–170. https://doi.org/10.1108/09590551311304338

Cardew, H. (n.d.). *The complete guide to DIY product photography*.

Cheng, F. F., Wu, C. S., & Yen, D. C. (2009). The effect of online store atmosphere on consumer's emotional responses - An experimental study of music and colour. *Behaviour and Information Technology*, *28*(4), 323–334. https://doi.org/10.1080/01449290701770574

Choi, J., & Nazareth, D. L. (2014). Repairing trust in an e-commerce and security context: An agent-based modeling approach. *Information Management and Computer Security*, *22*(5), 490–512. https://doi.org/10.1108/IMCS-09-2013-0069

Ciotti, G. (2013). *7 Marketing Lessons from Eye-Tracking Studies*. 1–13. Retrieved from https://neilpatel.com/blog/eye-tracking-studies/

Ding, C. G., & Lin, C. H. (2012). How does background music tempo work for online shopping? *Electronic Commerce Research and Applications*, *11*(3), 299–307. https://doi.org/10.1016/j.elerap.2011.10.002

Hsin Chang, H., & Wang, H. W. (2011). The moderating effect of customer perceived value on online shopping behaviour. In *Online Information Review* (Vol. 35). https://doi.org/10.1108/14684521111151414

Jain, S., Khan, M. N., & Mishra, S. (2017). Understanding consumer behavior regarding luxury fashion goods in India based on the theory of planned behavior. *Journal of*

Asia Business Studies, *11*(1), 4–21. https://doi.org/10.1108/JABS-08-2015-0118

Janssen, C., Vanhamme, J., Lindgreen, A., & Lefebvre, C. (2014). The Catch-22 of Responsible Luxury: Effects of Luxury Product Characteristics on Consumers' Perception of Fit with Corporate Social Responsibility. *Journal of Business Ethics*, *119*(1), 45–57. https://doi.org/10.1007/s10551-013-1621-6

Khan, U. (2018). *How to Use Images / Visuals to Boost W Conversion Rate How to Use Images to Boost Conversion Rat*. 1–20.

Kim, L. (2015). 16 Eye-Popping Statistics You Need to Know About Visual Content Marketing | Inc.com. *Inc.*, 16–19. Retrieved from https://www.inc.com/larry-kim/visual-content-marketing-16-eye-popping-statistics-you-need-to-know.html

Kim, Y. (2018). Power moderates the impact of desire for exclusivity on luxury experiential consumption. *Psychology and Marketing*, *35*(4), 283–293. https://doi.org/10.1002/mar.21086

Kohavi, R., & Thomke, S. (2017). The Surprising Power of Online Experiments. *Harvard Business Review*, *1617*(October), 1–15. https://doi.org/10.1126/science.1166632

Kuo, Y. F., & Wu, C. M. (2012). Satisfaction and post-purchase intentions with service recovery of online shopping websites: Perspectives on perceived justice and emotions. *International Journal of Information Management*, *32*(2), 127–138. https://doi.org/10.1016/j.ijinfomgt.2011.09.001

Li, H., Sarathy, R., & Zhang, J. (2014). The Role of Emotions in Shaping Consumers' Privacy Beliefs about Unfamiliar Online Vendors. *Journal of Information Privacy and Security*, *4*(3), 36–62. https://doi.org/10.1080/2333696x.2008.10855845

Lievonen, M. (2018). *Consumer Emotions and E-commerce: A Literature Review*. 385–402. https://doi.org/10.18690/978-961-286-043-1.27

Lynn, M., & Harris, J. (1997). Individual differences in the pursuit of self-uniqueness through consumption. *Journal of Applied Social Psychology*, *27*(21), 1861–1883. https://doi.org/10.1111/j.1559-1816.1997.tb01629.x

McKinsey & Company. (2016). The age of digital enlightenment. *McKinsey & Company*, 279–286. Retrieved from http://www.uk.logicalis.com/globalassets/united-kingdom/microsites/real-time-generation/realtime-generation-2016-report.pdf

Pegram, B. (2008). *Posing Techniques*.

Peiris, P., Kulkarni, D., & Mawatha, C. (2015). Implications of trust and usability on e-commerce adoption. *International Journal of Business and Information*, *10*(4), 519–557.

Simon, S. D. (2018). *How Keep Up With Trends: 10 Ways to Remain Competitive*. 1–9. Retrieved from https://www.yokellocal.com/blog/how-to-keep-up-with-trends-in-you-industry

Snyder, C.R. and Fromkin, H. L. (1980). Uniqueness: The Human Pursuit of Difference. *Plenum Press, New York*, *145*(10.1007/978-1-4684-3659-4).

Turban, E., Outland, J., King, D., Lee, J. K., Liang, T.-P., & Turban, D. C. (2018). *Electronic Commerce 2018*. https://doi.org/10.1007/978-3-319-58715-8

Wang, Q., Yang, Y., Wang, Q., & Ma, Q. (2014). The effect of human image in B2C website design: An eye-tracking study. *Enterprise Information Systems*, *8*(5), 582–

605. https://doi.org/10.1080/17517575.2014.925585

Wu, C. S., Cheng, F. F., & Yen, D. C. (2008). The atmospheric factors of online storefront environment design: An empirical experiment in Taiwan. *Information and Management*, *45*(7), 493–498. https://doi.org/10.1016/j.im.2008.07.004

APPENDIX A – Product photography sweep

APPENDIX B – Add to Wishlist Rate Uplift

Global					
	Uplift	**Diff**	**SE**	**z-score**	**p value**
Alternative	-0.7420%	-0.0389%	0.0485%	-0.8016	0.4226
By Visitor Type					
New	**Uplift**	**Diff**	**SE**	**z-score**	**p value**
Alternative	3.2463%	0.0674%	0.0597%	1.1277	0.2594
Returning	**Uplift**	**Diff**	**SE**	**z-score**	**p value**
Alternative	-1.1851%	-0.0764%	0.0627%	-1.2172	0.2236
By Device					
Desktop	**Uplift**	**Diff**	**SE**	**z-score**	**p value**
Alternative	-4.5939%	-0.2735%	0.0746%	-3.6666	0.0002***
Mobile	**Uplift**	**Diff**	**SE**	**z-score**	**p value**
Alternative	4.4075%	0.1815%	0.0632%	2.8741	0.0041**
Tablet	**Uplift**	**Diff**	**SE**	**z-score**	**p value**
Alternative	-0.4776%	-0.0449%	0.2831%	-0.1586	0.8377
By Country					
Australia	**Uplift**	**Diff**	**SE**	**z-score**	**p value**
Alternative	0.5196%	0.0319%	0.2703%	0.1179	0.9061
Brazil	**Uplift**	**Diff**	**SE**	**z-score**	**p value**
Alternative	4.6215%	0.1829%	0.1908%	0.9582	0.338
China	**Uplift**	**Diff**	**SE**	**z-score**	**p value**
Alternative	-10.6447%	-0.2943%	0.1319%	-2.2312	0.0257*
France	**Uplift**	**Diff**	**SE**	**z-score**	**p value**
Alternative	-15.8643%	-0.7013%	0.2056%	-3.4117	0.0006***
Germany	**Uplift**	**Diff**	**SE**	**z-score**	**p value**
Alternative	4.7488%	0.2278%	0.1845%	1.2343	0.2171
HK	**Uplift**	**Diff**	**SE**	**z-score**	**p value**
Alternative	4.2433%	0.2308%	0.3622%	0.6370	0.5241
Italy	**Uplift**	**Diff**	**SE**	**z-score**	**p value**
Alternative	4.0017%	0.1589%	0.2086%	0.7621	0.446
Japan	**Uplift**	**Diff**	**SE**	**z-score**	**p value**
Alternative	1.8310%	0.1240%	0.2050%	0.6049	0.5452
Russia	**Uplift**	**Diff**	**SE**	**z-score**	**p value**
Alternative	1.9449%	0.1851%	0.2228%	0.8311	0.406
UK	**Uplift**	**Diff**	**SE**	**z-score**	**p value**
Alternative	-9.8819%	-0.3370%	0.1265%	-2.6642	0.0077**
US	**Uplift**	**Diff**	**SE**	**z-score**	**p value**
Alternative	-9.9993%	-0.5219%	0.1362%	-3.8311	0.0001***

Note: *, ** and *** indicate statistical significance at a 95%, 99% and 99.9% level, respectively

APPENDIX C – Add to Bag Rate Uplift

Global					
	Uplift	**Diff**	**SE**	**z-score**	**p value**
Alternative	0.2772%	0.0247%	0.0622%	0.3976	0.6909
By Visitor Type					
New	**Uplift**	**Diff**	**SE**	**z-score**	**p value**
Alternative	0.5183%	0.0307%	0.0982%	0.3121	0.7550
Returning	**Uplift**	**Diff**	**SE**	**z-score**	**p value**
Alternative	0.2513%	0.0253%	0.0771%	0.3279	0.7430
By Device					
Desktop	**Uplift**	**Diff**	**SE**	**z-score**	**p value**
Alternative	0.3215%	0.0283%	0.0904%	0.3133	0.7541
Mobile	**Uplift**	**Diff**	**SE**	**z-score**	**p value**
Alternative	-0.2050%	-0.0181%	0.0892%	-0.2028	0.8391
Tablet	**Uplift**	**Diff**	**SE**	**z-score**	**p value**
Alternative	3.5026%	0.3853%	0.3061%	1.2585	0.2082
By Country					
Australia	**Uplift**	**Diff**	**SE**	**z-score**	**p value**
Alternative	-0.8166%	-0.1009%	0.3696%	-0.2729	0.7849
Brazil	**Uplift**	**Diff**	**SE**	**z-score**	**p value**
Alternative	8.6750%	0.4567%	0.2206%	2.0699	0.0385*
China	**Uplift**	**Diff**	**SE**	**z-score**	**p value**
Alternative	-5.9540%	-0.1955%	0.1451%	-1.3475	0.1777
France	**Uplift**	**Diff**	**SE**	**z-score**	**p value**
Alternative	1.4984%	0.1263%	0.2902%	0.4352	0.6634
Germany	**Uplift**	**Diff**	**SE**	**z-score**	**p value**
Alternative	-2.6995%	-0.2759%	0.2571%	-1.0731	0.2833
HK	**Uplift**	**Diff**	**SE**	**z-score**	**p value**
Alternative	8.0081%	0.8609%	0.4984%	1.7274	0.0841
Italy	**Uplift**	**Diff**	**SE**	**z-score**	**p value**
Alternative	-3.2307%	-0.2197%	0.2642%	-0.8318	0.4054
Japan	**Uplift**	**Diff**	**SE**	**z-score**	**p value**
Alternative	1.4475%	0.0762%	0.1820%	0.4185	0.6756
Russia	**Uplift**	**Diff**	**SE**	**z-score**	**p value**
Alternative	-8.4414%	-0.8402%	0.2219%	-3.7859	0.0002***
UK	**Uplift**	**Diff**	**SE**	**z-score**	**p value**
Alternative	0.9733%	0.0877%	0.2050%	0.4281	0.6686
US	**Uplift**	**Diff**	**SE**	**z-score**	**p value**
Alternative	-0.5229%	-0.0718%	0.2157%	-0.3328	0.7391

Note: *, ** and *** indicate statistical significance at a 95%, 99% and 99.9% level, respectively

APPENDIX D – Started Checkout Rate Uplift

Global					
	Uplift	**Diff**	**SE**	**z-score**	**p value**
Alternative	-0.1189%	-0.0034%	0.0362%	-0.0931	0.9259
By Visitor Type					
New	**Uplift**	**Diff**	**SE**	**z-score**	**p value**
Alternative	2.0396%	0.0394%	0.0575%	0.6854	0.4921
Returning	**Uplift**	**Diff**	**SE**	**z-score**	**p value**
Alternative	-0.5916%	-0.0188%	0.0449%	-0.4190	0.6752
By Device					
Desktop	**Uplift**	**Diff**	**SE**	**z-score**	**p value**
Alternative	-0.5586%	-0.0177%	0.0558%	-0.3178	0.7505
Mobile	**Uplift**	**Diff**	**SE**	**z-score**	**p value**
Alternative	-1.7833%	-0.0449%	0.0490%	-0.9147	0.3602
Tablet	**Uplift**	**Diff**	**SE**	**z-score**	**p value**
Alternative	17.9804%	0.5001%	0.1665%	3.0042	0.0027**
By Country					
Australia	**Uplift**	**Diff**	**SE**	**z-score**	**p value**
Alternative	-0.1018%	-0.0050%	0.2439%	-0.0207	0.9832
Brazil	**Uplift**	**Diff**	**SE**	**z-score**	**p value**
Alternative	-8.4641%	-0.1555%	0.1273%	-1.2218	0.2217
China	**Uplift**	**Diff**	**SE**	**z-score**	**p value**
Alternative	-3.4899%	-0.0344%	0.0809%	-0.4253	0.6708
France	**Uplift**	**Diff**	**SE**	**z-score**	**p value**
Alternative	11.8877%	0.2749%	0.1609%	1.7079	0.0877
Germany	**Uplift**	**Diff**	**SE**	**z-score**	**p value**
Alternative	-4.5186%	-0.1595%	0.1559%	-1.0235	0.3058
HK	**Uplift**	**Diff**	**SE**	**z-score**	**p value**
Alternative	0.9725%	0.0382%	0.3080%	0.1240	0.9013
Italy	**Uplift**	**Diff**	**SE**	**z-score**	**p value**
Alternative	-4.6295%	-0.0796%	0.1358%	-0.5860	0.5579
Japan	**Uplift**	**Diff**	**SE**	**z-score**	**p value**
Alternative	2.4860%	0.0489%	0.1136%	0.4310	0.6665
Russia	**Uplift**	**Diff**	**SE**	**z-score**	**p value**
Alternative	-7.4187%	-0.1531%	0.1055%	-1.4518	0.1465
UK	**Uplift**	**Diff**	**SE**	**z-score**	**p value**
Alternative	-4.1941%	-0.1551%	0.1334%	-1.1625	0.2448
US	**Uplift**	**Diff**	**SE**	**z-score**	**p value**
Alternative	3.1405%	0.1354%	0.1284%	1.0544	0.2917

Note: *, ** and *** indicate statistical significance at a 95%, 99% and 99.9% level, respectively

APPENDIX E – Conversion Rate Uplift

Global					
	Uplift	**Diff**	**SE**	**z-score**	**p value**
Alternative	1.3274%	0.0163%	0.0241%	0.6752	0.4995
By Visitor Type					
New	**Uplift**	**Diff**	**SE**	**z-score**	**p value**
Alternative	-1.9095%	-0.0138%	0.0350%	-0.3932	0.6943
Returning	**Uplift**	**Diff**	**SE**	**z-score**	**p value**
Alternative	1.9879%	0.0281%	0.0304%	0.9250	0.3550
By Device					
Desktop	**Uplift**	**Diff**	**SE**	**z-score**	**p value**
Alternative	0.2660%	0.0041%	0.0392%	0.1042	0.9170
Mobile	**Uplift**	**Diff**	**SE**	**z-score**	**p value**
Alternative	0.2066%	0.0018%	0.0296%	0.0622	0.9504
Tablet	**Uplift**	**Diff**	**SE**	**z-score**	**p value**
Alternative	15.8561%	0.2410%	0.1234%	1.9541	0.0507
By Country					
Australia	**Uplift**	**Diff**	**SE**	**z-score**	**p value**
Alternative	-5.5685%	-0.1347%	0.1705%	-0.7900	0.4295
Brazil	**Uplift**	**Diff**	**SE**	**z-score**	**p value**
Alternative	-1.3891%	-0.0094%	0.0791%	-0.1189	0.9053
China	**Uplift**	**Diff**	**SE**	**z-score**	**p value**
Alternative	38.0011%	0.1193%	0.0504%	2.3674	0.0179*
France	**Uplift**	**Diff**	**SE**	**z-score**	**p value**
Alternative	-5.8103%	-0.0614%	0.1049%	-0.5855	0.5579
Germany	**Uplift**	**Diff**	**SE**	**z-score**	**p value**
Alternative	-4.9483%	-0.0855%	0.1099%	-0.7777	0.4366
HK	**Uplift**	**Diff**	**SE**	**z-score**	**p value**
Alternative	20.5168%	0.3048%	0.2007%	1.5185	0.1289
Italy	**Uplift**	**Diff**	**SE**	**z-score**	**p value**
Alternative	-2.6046%	-0.0164%	0.0831%	-0.1975	0.8430
Japan	**Uplift**	**Diff**	**SE**	**z-score**	**p value**
Alternative	1.6039%	0.0138%	0.0752%	0.1829	0.8549
Russia	**Uplift**	**Diff**	**SE**	**z-score**	**p value**
Alternative	-5.6468%	-0.0566%	0.0742%	-0.7625	0.4455
UK	**Uplift**	**Diff**	**SE**	**z-score**	**p value**
Alternative	-6.9759%	-0.1025%	0.0844%	-1.2138	0.2247
US	**Uplift**	**Diff**	**SE**	**z-score**	**p value**
Alternative	7.9757%	0.1647%	0.0909%	1.8109	0.0702

Note: *, ** and *** indicate statistical significance at a 95%, 99% and 99.9% level, respectively

APPENDIX F – Return Rate Uplift

Global					
	Uplift	Diff	SE	z-score	p value
Alternative	1.2483%	0.2551%	0.7751%	0.3291	0.75
By Visitor Type					
New	**Uplift**	**Diff**	**SE**	**z-score**	**p value**
Alternative	-21.1202%	-3.8206%	1.7529%	-2.1796	0.0294*
Returning	**Uplift**	**Diff**	**SE**	**z-score**	**p value**
Alternative	4.8177%	1.0066%	0.8590%	1.1719	0.2412
By Device					
Desktop	**Uplift**	**Diff**	**SE**	**z-score**	**p value**
Alternative	0.4268%	0.0974%	1.0511%	0.0927	0.9261
Mobile	**Uplift**	**Diff**	**SE**	**z-score**	**p value**
Alternative	3.6444%	0.5639%	1.1867%	0.4752	0.6346
Tablet	**Uplift**	**Diff**	**SE**	**z-score**	**p value**
Alternative	-2.8370%	-0.7261%	3.2495%	-0.2235	0.8228
By Country					
Australia	**Uplift**	**Diff**	**SE**	**z-score**	**p value**
Alternative	23.4755%	3.932%	2.7800%	1.4143	0.1573
Brazil	**Uplift**	**Diff**	**SE**	**z-score**	**p value**
Alternative	-	-	-	-	-
China	**Uplift**	**Diff**	**SE**	**z-score**	**p value**
Alternative	71.3992%	19.671%	6.1380%	3.2048	0.0014**
France	**Uplift**	**Diff**	**SE**	**z-score**	**p value**
Alternative	-37.7019%	-9.380%	3.9359%	-2.3833	0.0172*
Germany	**Uplift**	**Diff**	**SE**	**z-score**	**p value**
Alternative	9.2316%	2.670%	2.9824%	0.8952	0.3707
HK	**Uplift**	**Diff**	**SE**	**z-score**	**p value**
Alternative	23.5294%	6.349%	5.4664%	1.1615	0.2454
Italy	**Uplift**	**Diff**	**SE**	**z-score**	**p value**
Alternative	22.9710%	3.492%	4.9513%	0.7052	0.4807
Japan	**Uplift**	**Diff**	**SE**	**z-score**	**p value**
Alternative	-9.8581%	-1.497%	3.0434%	-0.4919	0.6227
Russia	**Uplift**	**Diff**	**SE**	**z-score**	**p value**
Alternative	-1.8905%	-0.334%	2.8886%	-0.1155	0.9077
UK	**Uplift**	**Diff**	**SE**	**z-score**	**Significance**
Alternative	7.8832%	1.659%	2.4411%	0.6796	0.4968
US	**Uplift**	**Diff**	**SE**	**z-score**	**Significance**
Alternative	-3.9858%	-0.950%	1.7667%	-0.5377	0.5906

Note: *, ** and *** indicate statistical significance at a 95%, 99% and 99.9% level, respectively