UNIVERSITY OF CALIFORNIA,
IRVINE

An Empirical Examination of Deceptive Counterfeiting Activities in Electronic Markets

DISSERTATION

submitted in partial satisfaction of the requirements
for the degree of

DOCTOR OF PHILOSOPHY

in Management

by

Ziyi Cao

Dissertation Committee:
Professor Sanjeev Dewan, Chair
Assistant Professor Tingting Nian
Assistant Professor Behnaz G. Bojd

2023

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

ACKNOWLEDGEMENTS

I would like to express the deepest appreciation to my committee chair, Professor Sanjeev Dewan, who has provided exceptional guidance and support throughout my six years' journey as a PhD student. His profound expertise and experiences navigated me through many challenges. His insightful perspectives and rigorous attitude pushed the boundaries of my research. His confidence and determination inspired me whenever I was down. Without his persistent help this dissertation would not have been possible. I am very grateful for having such a passionate and knowledgeable scholar as him to be my advisor and shape my academic development.

I would like to thank my committee members, Professor Tingting Nian and Professor Behnaz G. Bojd, who provided very insightful feedback and suggestions to this dissertation. They never hesitate to express confidence in my ability and acknowledgement of my efforts, which not only boosted my resilience at these most difficult moments, but also had an impact on my career philosophy as a future faculty member. I sincerely appreciate that Professor Nian offered me the opportunity to work with her and connected me to other outstanding researchers. Her attentive and responsive guidance refined my understanding of empirical research and enhanced the overall quality of my work.

I am thankful to Professor Mingdi Xin, Professor Vidyanand (VC) Choudhary, Professor Vibhanshu Abhishek, Professor Francis MacCrory, and Professor Vijay Gurbaxani. They provided invaluable feedback and advice for my qualification exam, advancement defense, and job application. I am thankful to Professor Jiawei Chen and Professor Imran Currim, who sat in my advancement committee and provided insightful suggestions on the design of my study. I appreciate the company and friendship of my PhD fellows. I enjoyed the collaboration we shared and the delightful time we spent together.

It is an amazing experience to join this program, come to this adorable city, make amazing friends, and spend six incredible years. I would like to thank Xiaoyi (Sylvia) Gao, Shengjun Mao, and Yi-Jen (Ian) Ho, who provided valuable suggestions and help for my job search. I would like to thank Shijia Wu, Xiaoqi Xu, Xuan Xie, Haonan Yin, Qi Yan, and Qiao Wang. The joyful time we spent together in the theaters, on the road, or at the bay, have enriched my life and constituted a most significant part of my memories of the west coast. I appreciate that Qi and Qiao shared every one of my important moments during the six years and made me feel never alone. I would like to thank my dearest friend Yadong Lu, for his company, understanding, and support along the journey. I am grateful for the adventure and growth we have shared. I hope our laughter will be as wild and endless as the pattern of Newport waves.

Last, I would like to express my deepest gratitude to my parents, for their unconditional love, support, and accepting me as who I am. They encouraged me to pursue the PhD degree, and anything that I wanted. Their belief in me has been a source of strength and seen me through thick and thin. Their optimism and open-mindedness have shaped my perspective and motivated me to reach new heights.

VITA

Ziyi Cao

2017	B.S. in Engineering,
	School of Information, Renmin University of China
	Beijing, China
2023	Ph.D. in Management
	Paul Merage School of Business, University of California, Irvine
	Irvine, US

FIELD OF STUDY

Platform Strategy, E-Commerce Platform, Online Counterfeit Goods, Sharing Economy

ABSTRACT OF THE DISSERTATION

An Empirical Examination of Deceptive Counterfeiting Activities in Electronic Markets

by

Ziyi Cao

Doctor of Philosophy in Management

University of California, Irvine, 2023

Professor Sanjeev Dewan, Chair

With the proliferation of third-party sellers, counterfeiting has become a serious source of friction in online marketplaces. Different from traditional counterfeiters who target consumers who consciously seek for cheap knockoffs, sellers of deceptive counterfeit products target the whole population of online shoppers and make their products indistinguishable from genuine products. I develop two identification approaches to identify deceptive counterfeit products in online marketplace. The first applies natural language processing techniques to Amazon product reviews to generate a listing-level counterfeit probability, which in turn is used to classify ASIN (Amazon Standard Identification Number) listings as likely counterfeit or likely authentic. The second leverages an AI-based anti-counterfeit project launched by Amazon as an exogenous shock to directly measure the likelihood of a listing being counterfeit. I focus on two product categories, one a taste-based experience good (men's fragrances) and the other being a utilitarian product (wireless cell phone chargers). I embed the estimated counterfeit probability into a BLP-type choice model to model consumers' decision-making process and investigate how counterfeiting intensity affects user demand and platform revenues. I confirm that consumer disutility is increasing in the counterfeit probability, more so for high-end or popular products. I further find a substitution effect between likely counterfeit and likely authentic products: a 10% decrease in the price of a likely counterfeit product is associated with an average 0.0011% decrease in the market share of a likely authentic product. I leverage the structural parameter estimates to run a number of counterfactual experiments. These experiments suggest that protecting authentic sellers by simply banning all likely counterfeit listings would drastically reduce platform revenues. Instead, the deployment of counterfeit detection algorithms, and reporting the results to users, would align the interests of authentic sellers with the welfare of the platform. Overall, my analysis provides a robust empirical examination for identifying deceptive online counterfeiting and understanding its impact on the various stakeholders of an online retail platform.

CHAPTER 1

INTRODUCTION

1.1 Introduction – Identifications, Impacts, and Strategies

Counterfeit products account for over a half-trillion dollars of trade and are responsible for the loss of 750,000 jobs worldwide (Bressler and Bressler 2018). The problem is particularly acute in online marketplaces due to the inherent information asymmetries between buyers and sellers (Dewan and Hsu 2004, Dimoka et al. 2012, Kennedy 2020). According to a Wall Street Journal investigation (Berzon et al. 2019), thousands of items on Amazon were found to be deceptive or unsafe — even on Amazon Prime. Although Amazon has launched a series of anti-counterfeiting policies to deter knockoffs, the problem is exacerbated by the growing dominance of third-party sellers on Amazon, who account for over 55% of overall sales (Statista 2021). Counterfeiters have blended into the population of third-party sellers, and product recommendation algorithms often present a mix of genuine and counterfeit products to the online consumer. Indeed, it is very difficult to distinguish genuine products from fake and unauthorized replicas — deceptive counterfeit products — which are practically indistinguishable in terms of price, description, pictures, packaging and delivery terms (Kennedy 2020)[1]. This makes deceptive counterfeiting an insidious and costly problem online, but one that has not received much research attention.

[1] This is in contrast to non-deceptive counterfeit products, like fake luxury handbags, where the seller does not hide the fact that the product is a knockoff, and the consumer is willing to buy a fake presumably because of the lower cost. These are not the focus of this study, which looks at deceptive counterfeit products where the user cannot easily tell whether the product is genuine or fake.

Prior work has pointed out that counterfeiting harms a genuine producer's incentive to innovate and hurts its profits by taking away market share (Cho and Ahn 2010, Qian et al. 2015, Wang et al. 2018). On the contrary, some studies have shown that counterfeiting has a promotional effect for the original product and potentially expands overall market size and profits (Hui and Png 2003, Qian 2014). However, almost all of the prior work has focused on software piracy (Cho and Ahn 2010, Hui and Png 2003), or on the impact of non-deceptive counterfeit products where consumers consciously adopt cheaper fake versions (e.g., Qian 2014). In contrast, my focus is on deceptive counterfeit products in an online marketplace, where the overall impact of counterfeiting on: the demand for genuine products, consumer utility, and platform profits — can go both ways, positive or negative. A positive impact might be due to the fact that some consumers may not care if they have purchased a knockoff if it functions well and comes at a reasonable price. Further, the spread of knockoffs may have a promotional effect for producers of the genuine products, increasing awareness and market size. The entry of third-party sellers can amplify these effects, driving increased platform transactions and profits. On the negative side, consumers who receive defective knockoffs suffer disutility and monetary loss. At the same time, genuine brands see increased price competition, spillover harm to reputation and loss of sales. The decreased confidence of both consumers and brands will impede transactions, negatively impacting platform profits as well. Accordingly, the net impact of deceptive counterfeiting is an open empirical question — and one that serves to motivate this study.

In this work, I aim to identify and characterize counterfeit activity in online marketplaces, study its economic impact on the demand for genuine products, consumer utility, and platform welfare, and explore possible platform strategies in response to it. I focus on Amazon.com, which is the most dominant e-Commerce platform in the U.S., and much of the world. Since deceptive counterfeit products are virtually indistinguishable from genuine ones in terms of description, price, and product characteristics, a key challenge is how to identify them in the first place. In this regard, I develop two complementary approaches to predict or directly measure the intensity of online counterfeit activities. The first identification approach exploits the fact that the "proof of the pudding is in the eating"; i.e., after actually purchasing and consuming a product consumers have a reasonably good idea about whether the product they purchased was fake or not — and some of them go on to share their positive or negative experiences through the platform's review system. Even so, a significant complication for any counterfeit identification effort on Amazon is that multiple sellers on this platform can (and do) sell under the same ASIN; individual reviews do not identify which specific seller the product came from. Thus, from a user's perspective, the prevalence of reviews indicating fake activity influences their perception of the likelihood of encountering a fake product when purchasing from a given ASIN. I embrace these identification challenges and develop a probabilistic machine learning methodology to characterize the intensity of counterfeiting at the level of Amazon ASIN listings. The output of this methodology is the probability of encountering a counterfeit product when purchasing from an ASIN, which in turn allows us to classify "likely authentic" and "likely counterfeit" products.

The second identification approach leverages an exogenous event that happened to Amazon US online market. In order to mitigate the threat of counterfeit trades and protect authentic sellers' profits, Amazon launched an AI-based anti-counterfeiting service named Project Zero to help original brands fight against the threat of knockoffs. This service includes three parts: automated scanning and deleting fake offers; a self-removal tool that brands can use to remove fake offers by themselves; a serialization service that assigns a unique code to each unit of registered products. As this service is provided to more and more registered brands, many offers provided by third-party sellers are removed from existing ASIN listings. With Project Zero combining original brands' knowledge of their products and Amazon's artificial intelligence technology, the percentage drop in the number of optional offers caused by the service can reasonably measure the previous intensity of counterfeit activities of each listing. The standardized percentage drop in the number of optional offers is treated as the main direct measure in the second identification. I also provide an alternative measure generated from consumer insights into authenticity.

The data set that I use in this study consists of product and review data publicly available on Amazon, which I sampled through web-scraping. To that end, I crawled all the products listed under a given category (e.g., men's fragrances), to characterize the totality of the market for that category on Amazon, keeping track of prices and sales ranks on a daily basis. Historical review data and information on reviewer activity are also captured. For the first identification, I use natural language processing (NLP) techniques, applied to online reviews, to characterize the intensity of counterfeiting in different ASIN's within the product category. Specifically, I apply a semi-supervised hierarchical topic model on the review texts to generate the most frequently mentioned topics along with consumer attitudes on product authenticity, quality, shipping and customer service. The distribution of topics across the reviews along with other numeric product and review features are deployed in machine learning classification models to identify the probability of a counterfeit encounter when purchasing from an ASIN listing. For the second identification, I scrape the number of optional offers history data before and after Project Zero from an Amazon price tracking platform. I also collect sales and price history data previous to the launch of Project Zero for an unbiased econometric analysis based on the second identification.

With the likelihood of a counterfeit encounter probabilistically identified, or the intensity of counterfeit activities within listing exogenously measured, I adopt a discrete choice model, with random coefficients correlated to consumer preferences, to estimate the impact of counterfeiting on market shares. Then I design three counterfactual experiments to proxy three different platform strategies dealing with online counterfeit products. I simulate 100,000 consumers based on the choice model estimation and let them make optional decisions under the three scenarios. I calculate product sales and platform revenues under each strategy to explore the economic significance of detecting knockoffs.

The results suggest that a higher counterfeiting probability significantly reduces consumer mean utility gained from purchasing under an ASIN listing, and the magnitude of the disutility from counterfeiting is larger for best-selling and for expensive products. Estimated cross-price elasticities indicate that likely counterfeit products exhibit significant substitution effects with respect to likely authentic products. Estimation of consumer decision-making process and economic impacts under the machine-learning-based identification and exogenous-event-based identification are consistent. Specifically, with the first identification, a 10% decrease in the price of a likely counterfeit product is associated with a 0.0011% average decrease in the market share of each related likely authentic product. And the substitution effect is relatively stronger for expensive brands, which are more frequent targets of counterfeiters. I leverage the structural parameter estimates to conduct some counterfactual policy experiments. The experiments indicate that a proliferation of counterfeit activity on the platform would negatively impact the market share of authentic sellers, and also reduce platform revenues. However, banning all listings with counterfeit activity would protect authentic sellers, but at the cost of a steep decline in platform revenues. A more moderate approach for the platform would be to deploy counterfeit detection algorithms — along the lines of what we propose — and report the results to the consumers. This regime would align the interests of authentic sellers and platforms alike. I discuss managerial and policy implications of the results in the concluding chapter.

In what follows, Section 1.2 provides a review of the relevant literature. Chapter 2 introduces the empirical context first and then outlines methodologies for two identifications of counterfeit intensity in parallel, with data description and model-free evidence included respectively. Chapter 3 describes the choice-modeling framework and empirical results under two different identifications, as well as the application on a utilitarian product category. Chapter 4 presents the design, implementation, and results of several counterfactual experiments. Finally, Chapter 5 provides the discussion and conclusions.

1.2 Literature Review

One of the areas of prior work that is related to this research is that on the impact of piracy and counterfeit sales; much of this work is theoretical in nature. A case in point is Sundararajan (2004), which models the optimal strategy of pricing and technology protection in response to digital piracy. Looking at the counterfeiting of information goods, the analysis of Cho and Ahn (2010) suggests that the threat of counterfeiting will reduce firms' incentive to innovate, and to choose a lower quality level. Similar adverse effects have been investigated in the case of physical goods (Qian et al. 2015) and for the mobile app market (Wang et al. 2018). Qian et al. (2015) shows that the counterfeiting issue will push authentic brands to make adjustments in product quality, in order to differentiate with counterfeiters and attain a separating equilibrium, as counterfeit sellers take away profits in a pooling equilibrium. Want et al. (2018) provides empirical evidence that high quality copycats of mobile apps will negatively affect the demand for original apps.

At the same time, there is contrary evidence of potential positive impact of counterfeiting, in that illegal copies can increase market size for the original when taking network effects into account (Givon et al. 1995, Hui and Png 2003). In this vein, some analytical studies develop game-theoretic models to show that under strong network effects (Jain 2008), or in a monopoly market (Lahiri and Dey 2013), piracy of intellectual properties will strategically increase product quality and firm profits. Consistent with these theoretical conclusions, Lu et al. (2020) explores the movie industry and finds that post-release piracy increases WOM volume, which in turn can have a positive impact on revenues. Turning to the traditional retail scenario for physical goods, Qian (2014) conducted a quasi-natural experiment to examine advertising and substitution effects of counterfeits and found the advertising effect dominates for high-end products.

Generally, prior work suggests that the impact of counterfeiting is complicated and heterogeneous across markets. However, much of the previous research either focuses on digital goods, on luxury goods, or on offline retailing. The counterfeiting issue in the setting of an online marketplace remains largely unexplored, in part due to the difficulty of identifying knockoffs, and the lack of a ground truth. Given that most consumers intentionally purchase pirated information goods and counterfeit luxury goods, to save on price, the results from those settings do not readily translate to the case of deceptive online counterfeiting, which is the focus of this study.

It should be noted that the detection of fake products in online platforms has been a popular topic in computer science. Specifically, a variety of deep learning algorithms have been developed to detect traces of counterfeiting in textual information and/or images posted by the seller. Specific approaches include comparing images posted by others combined with seller information on social network platforms (Cheung et al. 2018), processing microscopic images of physical products (Sharma et al. 2017), or identifying matches for specific products using crawled information (Chaloux et al. 2020). However, these methods are either embedded in devices or utilize photos of physical items as inputs, which do not transfer into consumers' perception of counterfeits in online marketplaces. As discussed earlier, purveyors of deceptively counterfeit products online tend to thoroughly imitate authentic sellers' behavior by posting realistic pictures and descriptions, making it difficult to distinguish counterfeit from genuine products — short of purchasing and consuming the products.

Accordingly, I turn to NLP methods applied to user-generated review information to decipher product quality and authenticity. Plenty of research documents the significant effect of eWOM on sales in online retail (Chevalier and Mayzlin 2006), hotel booking (Lewis and Zervas 2016) and mediating consumer experience (Bai et al. 2020). In addition, textual information of reviews has also proven to have an economic impact. For example, Archak et al. (2011) applies NLP techniques to extract consumer opinions on product attributes from reviews and explore their impact on sales. Ghose et al. (2011) uses text mining to generate consumer preferences on hotel features and develops a hotel ranking system based on a choice model. Yet none of them have used the topic features extracted from reviews for identification of counterfeiting intensity in specific product markets. I propose an approach which leverages semi-supervised topic modeling to capture consumers' feedback on product authenticity, and then use the outputs from topic modeling as predictors to train machine learning classification models to identify counterfeiting activity.

CHAPTER 2

Identifications of Online Counterfeit Products

2.1 Empirical Context

Counterfeits are products bearing the trademark of a legitimate brand but not made by the original manufacturers. Although they are illegal copies, there have been some active markets for them for a very long time. Some industries where knockoffs are commonly seen include but not limited to luxury goods and information goods. In the recent decade, as e-commerce platforms emerge and grow exponentially, counterfeiters are also taking the advantage and infringe upon much more variety of product categories than ever. And proliferation of third-party sellers rooted in the nature of platform economy accelerates the explosion of online counterfeiting activities. Compared to traditional counterfeiting activities, counterfeit products in electronic markets can be very confusing due to the intrinsic information asymmetry in online shopping. As a result, deceptive counterfeit products are scaling up and targeting the whole population of online shoppers.

When deceptive counterfeiters set their product prices close to authentic products' prices and design their product pages as attractive as genuine ones', with detailed information and images, consumers' search costs will greatly increase. While the review system can disclose some information and help consumers avoid risky products, not all consumers are equally sensitive or patient to online reviews. And consumers who unconsciously buy knockoffs suffer from a monetary loss. On the seller side, the competition from counterfeit sellers is driving some original manufactures away from e-commerce platforms like Amazon. However, on the other hand, some brands, especially newcomers, can benefit from a promotional effect if counterfeit products are spreading on the market. For consumers who care more about the basic functionality instead of the excellence of quality or social values attached to a specific brand, they can also gain some positive utility from buying a knockoff. When the entry of third-party sellers amplifies these positive effects, platforms gain additional profits from increased transactions.

Due to the complication of the counterfeiting issue, for a long time, many giant platforms across the globe – such as Alibaba, Amazon, and Flipcart – used to only rely on consumers or sellers to report on suspicious listings but did not actively detect and remove them. As the threat of counterfeiting activities became more and more salient, Amazon recently took one step further to launch Project Zero, an AI-based subscription service, to confront the issue. It is a self-service tool that original brands can make use of to distinguish their products from knockoffs (Hernandez, 2021). When brands previously registered on Amazon enroll their products and upload trademarks, logos or other key data points, Amazon will automatically scan the marketplace to find and remove suspected fake products. At the same time, it empowers registered brands to directly remove illegal copycats from the store. Brands can also enroll products in a serialization service, which assigns a unique QR code as a genuine certification to each unit from manufacturers (Amazon, 2023). This project was initially launched in February 2019, to a limited number of original brands in the US market, and gradually expanded to more brands and other countries.

This study is built on Amazon, the most popular e-commerce platform. First of all, it has a very mature and active reputation system where users can post reviews with bare limitations for their verified purchases. Information of authenticity can be partly disclosed by historical reviews and perceived by future consumers. Second, the launch of Project Zero becomes an exogenous shock to the distribution of Amazon sellers and can be leveraged to measure the intensity of online counterfeiting activities. Last but not least, the public sales rank information on Amazon well presents the dynamic of market competition.

According to a survey asking consumers what the product categories they think mostly contaminated by counterfeit trades are, fragrances, cosmetics, sporting goods, and bags got the majority votes (LocalCircles, 2018). Another e-commerce giant Ebay also warns their users to be cautious of the riskiest categories including electronics, sneakers, toys, etc. (Ebay, 2023). Therefore, I choose to focus the main analysis on one of the categories known to be impacted by online counterfeiting – men's fragrances (see, e.g., Quora 2019, Steele 2019, Silcox 2021, Pieterse 2021). I also study the cell phone wireless charger category to extend the empirical framework to a utility product category.

The analysis is at the level of unique Amazon Standard Identification Number (ASIN) — i.e., the listing level. In other words, each item in the data sample corresponds to one listing on Amazon.com, identified by a unique ASIN. Note that multiple sellers may be selling the same product under the same ASIN, which is particularly common in the men's fragrances category. It should also be noted that sometimes the same product is sold under different ASINs, in which case I consider the different ASIN listings and associated suppliers to be independent.

In a scenario where there are multiple *optional offers* provided by different sellers under an ASIN, one of those sellers (primary) is placed by Amazon in the so-called Buy Box in the ASIN listing, whereas the other alternative sellers can be accessed from links on the side of the listing (see Figure 1). Which seller wins the Buy Box and becomes the featured seller is determined on a dynamic basis by Amazon, based on prices, shipment and other seller performance metrics such as the feedback rate, customer response time, and the like (Zeibak 2020). While the featured seller changes over time, and consumers can purchase the item from different sellers listed on the ASIN, the historical reviews are all aggregated under the same ASIN listing and there is no way to recover which review corresponded to which seller. Thus, when consumers see reviews indicating past fake product transactions, they can only draw inferences about the overall likelihood of encountering a fake product under an ASIN and are not able to resolve which of the sellers is likely to be shipping counterfeit products. Meanwhile, when Amazon Project Zero detects suspicious offers selling fake products, it removes these options from corresponding ASINs. Although sellers can list or withdraw their products from existing listings for different reasons, a significant drop in the number of *optional offers* after the launch of Project Zero can be considered as a good signal of counterfeiting activities within a specific ASIN. Accordingly, I use ASIN-level specified counterfeiting probability or intensity to capture the likelihood that a consumer would encounter a fake product when purchasing under a given ASIN.

Figure 1. A Typical Amazon ASIN Listing with Multiple Sellers

2.2 Identification I – Machine Learning Detection Based on User-Generated Reviews

2.2.1 Data and Variables

I scrape data on products in the category of men's fragrances and the category of cell phone wireless chargers on Amazon, and compile detailed product and review information, both numeric and textual. To ensure that the data set covers the entire product market, I collect all the products displayed in the search results, up to the last page. After manually removing a few items mistakenly included for they have similar matching names but do not belong to the focal category, I formed a refined sample and scraped their prices, sales ranks and other time-varying features on a daily basis from December 2020 to April 2021. Coupon and discount information are also obtained to adjust for the actual price paid by the consumer. Sales ranks are converted first to sales, and then into market shares — which are the dependent variables in the BLP-type empirical models, a la Berry et al. (1995), as I discuss in detail in Chapter 3.

The men's fragrance product category which is used in the main analysis contains 5661 products in 52 daily (aggregated to 10 weekly) periods. And the category of cell phone wireless chargers contains 1120 products in 120 daily (aggregated to 17 weekly) periods. Considering that the market is rather competitive with a long-tailed distribution of listing market shares, to achieve a better fit in the choice models, I use the 1037 highest-ranked products whose individual market shares in Amazon are greater than 0.01%; these add up to about 80% of the overall market. Summary statistics of sales (panel) data are shown in Table 1a. Similar processing is conducted on the cell-phone wireless charger category, and that sample contains 565 products in 17 weekly periods.

Table 1a. Summary Statistics of Sales Data

Variables	Description	(1) N	(2) mean	(3) S.D.	(4) min	(5) max
Share	Market Share (%)	9,520	0.0221	0.231	9.12e-06	8.67
Price	Unit Price of Item ($)	9,520	42.12	32.51	2.560	380.8
Sales	Estimated Sales of Item (count)	9,520	35.73	381.6	0.0131	14,679
Num_Reviews	Number of User Reviews	9,520	874.5	2,188	5.667	48,266
Amazon's Choice	Dummy Variable	9,520	0.222	0.383	0	1
Rating	Average Weekly Rating Valence	9,520	4.561	0.188	3	5
Rank	Product Sales Rank	9,520	88,338	58,905	199.8	522,204

Notes: Sales data is based on the weekly average amount. Num_Reviews and Rating are cumulative values in the current week.

Finally, I collect all historical reviews of each listing in the sample, including: the overall rating; the number of helpful votes that a review received; the number of pictures shared under a review; and the textual comments. I also collect details such as whether the review is linked with a verified purchase, a Vine voice, or posted by a "top reviewer." I applied NLP techniques on review texts to extract topics related to product quality and authenticity, which are key inputs to the machine learning models to estimate counterfeiting probability

— see Table 1b for summary statistics of the product data for the sample of men's fragrances.

Table 1b. Summary Statistics of Product Data

	Variable	Description	(1) N	(2) mean	(3) S.D.	(4) min	(5) max
Product Characteristics	Size	Perfume size in oz	1,037	3.628	1.800	0.0300	17
	Vol 1	= 1 if Eau de Cologne	1,037	0.217	0.412	0	1
	Vol 2	= 1 if Eau de Toilette	1,037	0.669	0.471	0	1
	Vol 3	= 1 if Eau de Parfum	1,037	0.103	0.304	0	1
	Vol 4	= 1 if Parfum	1,037	0.00675	0.0819	0	1
	Shipping	Categorical variable	1,037	2.777	0.493	0	3
	Free Return	= 1 if free return	1,037	0.00771	0.0875	0	1
	Num_Options	# sellers listed per ASIN	1,037	13.96	12.10	1	63
	Price S.D.	Std Dev of Price in ASIN	1,037	6.334	6.245	0	86.03
	Num_Disclosing	# disclosing reviews	1,037	5.878	16.27	0	209
	Disclosing Ratio	% of disclosing reviews	1,037	0.0257	0.0354	0	0.300
	Counterfeit_10_01	=1 if Num_Disclosing>=10 and Disclosing Ratio<0.01	1,037	0.135	0.342	0	1
	Counterfeit_5_01	= 1 if Num_Disclosing>=5 and Disclosing Ratio >0.01	1,037	0.236	0.425	0	1
Reviews	Review_1_star_pct	% of one-star reviews	1,037	0.0407	0.0270	0	0.200
	Review_3_star_pct	% of three-star reviews	1,037	0.0502	0.0271	0	0.240
	Review_5_star_pct	% of five-star reviews	1,037	0.782	0.0757	0.410	0.950
	Num_Reviews_Text	Number of reviews with text	1,037	197.9	424.9	1	5,000
Review Metrics	Num_Helpful_Votes	No. helpful votes /r	1,037	0.733	1.615	0.0357	48.47
	Num_Images	No. images /r	1,037	0.0369	0.0787	0	1.400
	Summary Wordcount	No. words in summary /r	1,037	3.360	0.593	1.333	7
(/r = Average per review)	Text Wordcount	No. words in text /r	1,037	19.67	8.788	4.143	96.18
	Overall	Text reviews avg rating	1,037	4.307	0.335	3	5
	Vine	% of Vine voice reviews	1,037	0.00376	0.0318	0	0.579
Topic Variables	Fake		1,037	0.0421	0.0488	0	0.375
	Like		1,037	0.145	0.0738	0	0.600
	Dislike		1,037	0.0430	0.0409	0	0.333
	Price_positive		1,037	0.0865	0.0621	0	0.667
	Price_negative	% of reviews	1,037	0.0495	0.0463	0	0.500
	Scent_positive	with this topic	1,037	0.0686	0.0507	0	0.400
	Scent_negative		1,037	0.0581	0.0452	0	0.300
	Last_positive		1,037	0.121	0.0707	0	0.667
	Last_negative		1,037	0.488	0.138	0	1
	Package_positive		1,037	0.00855	0.0147	0	0.143
	Shipping_positive		1,037	0.0240	0.0299	0	0.231
Classification Results	Counterfeit (0/1)	= 1 if likely counterfeit	1,037	0.335	0.472	0	1
	Counterfeit Prob	Probability of Counterfeit	1,037	0.424	0.233	0.01000	0.990

Notes: Product data for the men's fragrances category. Disclosing reviews are ones that explicitly identify fake products, using terms like "fake," "counterfeit" or "knockoff." Review Metrics and Topic Variables are within-product average of review-level features.

2.2.2 Model-Free Evidence

Before conducting NLP techniques, I first use keyword matching to identify the reviews that specifically complained about having bought a fake product, so that we can gain an initial understanding of topic distributions and the potential of reviews to predict the probability of counterfeiting. In particular, I label the texts containing keywords like "fake", "counterfeit", "knockoff", "not real", etc. as disclosing reviews and examine the number and percentage of such reviews in each listing. The distribution of disclosing review count and ratio are plotted in Figure 2, from which we can tell that most products have less than 30 disclosing reviews, that constitute less than 5% of all historical reviews. Although disclosing information is rather sparse in the review data, they serve as good indicators of suspicious knockoffs since a large proportion of listings never receive any disclosing reviews at all.

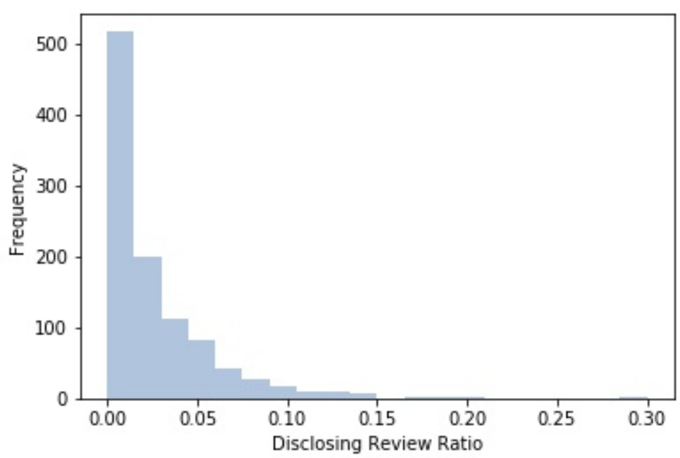

Figure 2. Percentage of Disclosing Reviews

After extracting topic variables from review texts using NLP techniques and manually labeling the training set, we are left with a correlation matrix and rating distribution as shown in Table 2. The labeled Fake Dummy is highly correlated with: the number of disclosing reviews; percentage of one-star and three-star ratings; and topic variables such as “fake”, “positivity” in “scent”, “lasting power”, and “price”.

Table 2. Correlation Matrix

	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16
1. Num_Reviews	1.000															
	(0.000)															
2. Num_Disclosing	0.766	1.000														
	(0.000)	(0.000)														
3. Disclosing Ratio	0.059	0.258	1.000													
	(0.440)	(0.001)	(0.000)													
4. Counterfeit_10_01	0.511	0.626	0.345	1.000												
	(0.000)	(0.000)	(0.000)	(0.000)												
5. Review_1_star_pct	-0.030	-0.149	-0.095	-0.170	1.000											
	(0.696)	(0.050)	(0.211)	(0.025)	(0.000)											
6. Review_3_star_pct	0.044	-0.151	-0.204	-0.202	0.620	1.000										
	(0.564)	(0.047)	(0.007)	(0.008)	(0.000)	(0.000)										
7. Review_4_star_pct	-0.037	-0.165	-0.233	-0.245	0.111	0.415	1.000									
	(0.624)	(0.030)	(0.002)	(0.001)	(0.145)	(0.000)	(0.000)									
8. Vine	-0.023	-0.055	-0.088	-0.082	0.101	0.223	0.128	1.000								
	(0.764)	(0.471)	(0.248)	(0.278)	(0.188)	(0.003)	(0.094)	(0.000)								
9. Fake	0.034	0.208	0.792	0.276	-0.099	-0.224	-0.232	-0.095	1.000							
	(0.656)	(0.006)	(0.000)	(0.000)	(0.193)	(0.003)	(0.002)	(0.208)	(0.000)							
10. Price_positive	0.007	0.028	-0.065	0.035	-0.056	0.030	-0.011	0.035	-0.049	1.000						
	(0.926)	(0.709)	(0.395)	(0.644)	(0.466)	(0.691)	(0.886)	(0.643)	(0.523)	(0.000)						
11. Scent_positive	-0.008	-0.055	-0.146	-0.091	0.180	0.230	-0.088	0.072	-0.140	-0.139	1.000					
	(0.913)	(0.469)	(0.053)	(0.228)	(0.018)	(0.002)	(0.251)	(0.344)	(0.063)	(0.066)	(0.000)					
12. Fake_h	0.198	0.431	0.516	0.511	-0.128	-0.259	-0.323	-0.087	0.675	0.014	-0.154	1.000				
	(0.010)	(0.000)	(0.000)	(0.000)	(0.100)	(0.001)	(0.000)	(0.260)	(0.000)	(0.852)	(0.045)	(0.000)				
13. Dislike_h	0.137	0.158	0.155	0.160	0.252	0.072	-0.131	0.042	0.251	-0.009	-0.100	0.405	1.000			
	(0.075)	(0.039)	(0.044)	(0.037)	(0.001)	(0.356)	(0.092)	(0.591)	(0.001)	(0.903)	(0.194)	(0.000)	(0.000)			
14. Last_positive_h	0.048	0.070	-0.067	0.046	-0.132	0.060	0.130	0.123	-0.154	-0.021	-0.046	-0.039	-0.061	1.000		
	(0.532)	(0.361)	(0.384)	(0.554)	(0.089)	(0.442)	(0.095)	(0.110)	(0.045)	(0.782)	(0.551)	(0.614)	(0.428)	(0.000)		
15. Overall_h	-0.195	-0.286	-0.216	-0.363	-0.154	0.024	0.113	-0.008	-0.281	0.041	0.183	-0.544	-0.467	0.074	1.000	
	(0.011)	(0.000)	(0.005)	(0.000)	(0.047)	(0.754)	(0.145)	(0.914)	(0.000)	(0.593)	(0.017)	(0.000)	(0.000)	(0.337)	(0.000)	
16. Fake Dummy	0.277	0.368	0.222	0.387	0.042	-0.182	-0.271	-0.142	0.233	-0.106	-0.156	0.380	0.291	-0.174	-0.349	1.000
	(0.000)	(0.000)	(0.003)	(0.000)	(0.584)	(0.017)	(0.000)	(0.059)	(0.002)	(0.161)	(0.039)	(0.000)	(0.000)	(0.023)	(0.000)	(0.000)

Lastly, I highlight the variations of sales and prices for products that are likely genuine or likely counterfeit (based on a 50% predicted probability cutoff). As shown in Figure 3a, likely counterfeits are more likely to have a higher price, as compared to likely genuine products. The pattern is suggestive of the fact that counterfeiters target products with higher prices. In terms of sales, I find that for products with lower sales, the composition of likely authentic products is higher relative to likely counterfeits. Both likely counterfeit and likely authentic products follow long-tailed distributions of sales (Figure 3).

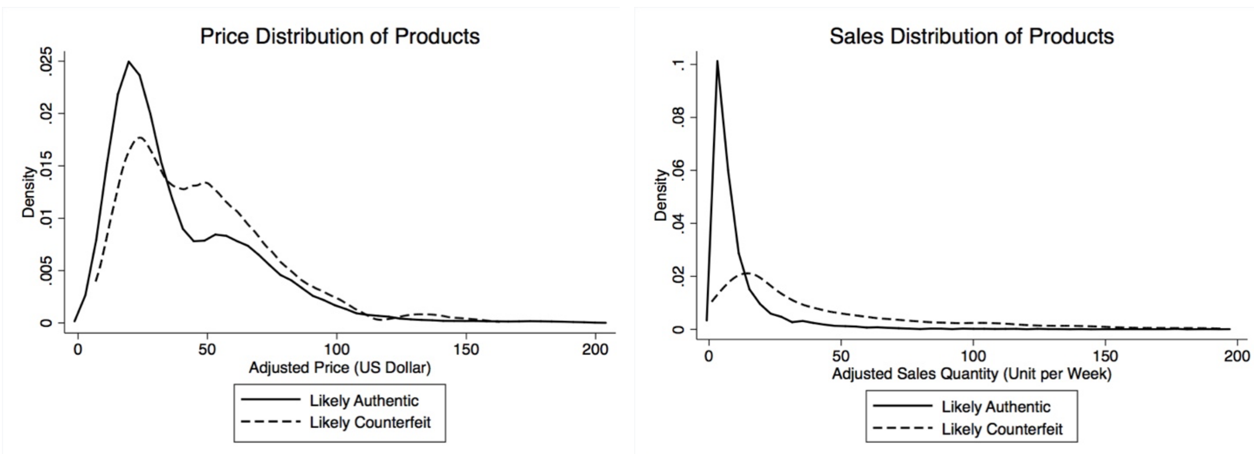

Figure 3. Price and Sales Distributions for Likely Authentic and Likely Counterfeit Products

2.2.3 Counterfeit Identification

As discussed earlier, the focus of this study is on deceptive counterfeit products, which are fake and unauthorized replicas of real goods, but are merchandised in a manner which makes them virtually indistinguishable from genuine products. Yet, the functionality of the fake products is generally inferior to that of the genuine ones. For example, in the category of fragrances, a listing of a fake product would have essentially an identical description of important features such as scents, weights, ingredients, country of origin, pictures and videos — like a genuine product. However, after consuming a fake perfume, consumers might notice that the scent seems strange or diluted, or fades rather fast. That is, the genuineness of the products, or lack thereof, can generally be discerned post-purchase, and this fact is likely to be revealed in user reviews — which is the key to the identification of counterfeits. Consumers are assumed to have a vertical preference for the authenticity of products; i.e., every consumer prefers a genuine product to a counterfeit one. And due to the existence of multiple sellers per listing, this study focuses on the probability of encountering a fake product when purchasing from a given ASIN listing.

The identification of counterfeit products in online markets has always been a challenge in empirical research due to the lack of ground truth. The most relevant and efficient method is contributed by Wang et al. (2018), who studied the mobile app market, and utilized text and image matching to cluster similar apps, and used launch dates to differentiate copycats from the original app (with the earliest launch date) within each cluster. However, similar methods cannot be applied to the e-Commerce scenario. First, on giant e-commerce platforms like Amazon, within each defined category there may exist thousands of related products; each product may be featured in multiple ASIN listings, and each listing may have multiple sellers, some of whom may be selling counterfeit products. Second, based on my definition of online deceptive counterfeit products, sellers will pretend and imitate the behavior and signals of authentic sellers in terms of pricing and posting product information; therefore, it is difficult to differentiate genuine products from knockoffs within each matched cluster simply by leveraging one feature such as price or launch date. Therefore, I turn to user-generated reviews and apply NLP techniques to extract frequently mentioned topics including product authenticity and other aspects of quality and use these generated indicators to train machine-learning classification models for prediction, as I describe next.

2.2.4 Topic Extraction

Our original sample of men's fragrance contains 5661 products, for which 352,933 reviews are collected up to April 2021. I pre-process the raw text by case normalization, tokenization (including removing punctuation), POS (part of speech) tagging and lemmatization according to POS tags. I also apply language detection to raw texts and keep only those written in English, which constitute over 85% of all reviews. Next, text vectorization is conducted to convert each pre-processed review text into a numeric vector. I choose the TF algorithm and keep the top 10,000 frequent tokens (words) in the vector, which is appropriate because review documents are relatively short and straightforward and contain no complicated semantic issues.

With textual data embedded into a numeric matrix, each row of which represents a piece of review, I train an anchored topic model on it to generate targeted topics and determine whether the documents contain corresponding information in each of the topics. Anchored CorEx (Correlation Explanation) is a semi-supervised hierarchical topic model which allows users to guide the topic generating direction by assigning anchored words to define topics of interest (Gallagher et al. 2017). The process can be described in two sequential stages: first, choosing the topics of interest (relevant to the counterfeit prediction in our case) and selecting anchored words; second, Model fitting and improvement. At the first stage, to decide which topics regarding product quality are most frequently mentioned by consumers either positively or negatively and find the anchored words to locate such topics, I fit an LDA model to obtain some independent product features with relevant keywords as a reference. For example, in the men's fragrance category, the most relevant topics generated by the unsupervised models are consumer sentiments, attitudes on scents (with keywords nice, pleasant, classy and cheap, overpower, weird, etc.), and attitudes on lasting power (with keywords long, throughout, all day or lost, flourish, etc.). Next, I randomly read a group of reviews to refine the construction of topics and anchor words, so that the topics of interest are better captured. Thirteen topics for men's fragrances related to counterfeit identification are defined as follows: explicit counterfeit (fake), overall sentiments, attitude towards price, scent, longevity, package and shipping (positive or negative). At the second stage, after fitting the Anchored CorEx model, I remove topics which are not accurately anchored or under which most of the predefined keywords were rarely captured; then I modify the model to include eleven topics in order to increase total correlation. I fit the refined model to extract indicators of texts under each topic and aggregate the review-level topic dummies into product-level percentage scores. For example, if one review states "This product is a cheap knock-off of the actual cologne. The box comes with white stickers on it to cover the actual serial info and the scent is off and does not last nearly as long", it will be tagged 1 under the topic of "Fake" and "Last_negative". The higher the product-level percentage scores under the topic of fake or other topics with negative sentiments towards a particular feature, the more likely the product will be deemed to be counterfeit.

For the purposes of further model improvement, I also calculate the topic dummies' average values weighted by the log of helpful votes each review receives (denoted as *Fake_w* for example), as well as the average of a subset including only reviews with at least one helpful vote (denoted as *Fake_h*). User-specified anchored words are consistent with model-determined keywords, as listed in Table 3. The review-level topic indicators are aggregated to product level as predictors in the classification of fake products.

Table 3. Anchored Topic Model

Topic	Anchored Words	Model Generated Keywords
Fake	Fake, defective, knock, knockoff, counterfeit, diluted, water	Fake, water, knock, counterfeit, knockoff, defective, real, cool
Like	Love, like, amazing, best, great, awesome, satisfy, favorite, complement	Like, favorite, compliment, best, great, lot, love, satisfy
Dislike	Waste, disappoint, bad, dislike, poor, terrible	Bad, waste, disappoint, terrible, poor, dislike, money, total
Price_positive	Price, deal, sale, value, worth, bargain, good	Good, price, value, worth, deal, sale, bargain, size
Price_negative	Critique, expensive	Expensive, di, gio, aqua, creed, aventus, acqua, similar
Scent_positive	Scent, attract, crisp, nice, classy, pleasant, delicious, classic	Scent, nice, pleasant, classic, delicious, crisp, classy, attract
Scent_negative	Overpower, strange, strong, cheap, weird, disgust, much	Strong, cheap, overpower, weird, strange, disgust, offensively, cheerful
Last_positive	Last, long, throughout, all, day	Long, day, time, stay, doesn, lasting, father, valentine
Last_negative	Minute, away, lost, flourish	Minute, away, cologne, bottle, just, say, try, use
Package_positive	Package, fancy	Package, fancy, et, est, je, le, tr, pa
Shipping_positive	Shipping, fast	Fast, shipping, delivery, ship, super, described, service, quick

Notes: CorEx (Anchored Correlation Explanation) model was used to generate these results. Anchored Words are the tokens used to describe the topics of interest for model fitting. Model Generated Words are the tokens extracted by the model to define the topics as part of the model output.

2.2.5 Classification Model for Likely Counterfeit Listings

As previously discussed, this work relies on Amazon review data to predict the likelihood of encountering a fake product when purchasing from a particular ASIN. One major challenge is that the review system is at the listing level. Individual reviews cannot be linked to specific sellers, so consumers have to draw imprecise inferences at the ASIN level, about the likelihood of experiencing a counterfeit product. I build a supervised classification model for predicting the probability of encountering a fake product. I manually constructed a training data set, wherein multiple human coders (graduate students in our case) browsed ASIN page information, rating distribution and textual reviews for all the reviews under a given ASIN, and generated a label for that ASIN, as either "Likely Authentic" or "Likely Counterfeit." The interpretation of the Likely Counterfeit (Likely Authentic, respectively) label is that there is a better than even perceived likelihood that a product purchased from this ASIN is going to be counterfeit (authentic, respectively). The labels were added to the data set by two coders independently agreeing on the label for each ASIN. ASINs for which there was no consensus among the coders were left out of the training data set.

Our approach to constructing the training data set is a departure from the traditional notion of a machine learning data set, where each instance has an objectively certain outcome, whereas in our case, the labels are the result of a subjective assessment. A number of factors contribute to an assessment that an ASIN should be coded "Likely Counterfeit." For one thing, an unusually high number of reviews that explicitly call the product fake — using some combination of keywords like "fake", "counterfeit", "knockoff" — were a sign of counterfeiting activity in the ASIN. Also, the coders looked at the distribution of ratings, and specifically, the proportion of 1-star ratings. Cases where there were more 1-star than 2-star ratings were examined more closely for counterfeiting activity. At the current stage, our training set for each category (men's fragrances and cell-phone chargers) contains 200 listings each.

To reduce possible bias caused by topic modeling, I also leverage the metrics of "disclosing" reviews as supports, if it explicitly complains about having bought a fake product and contains one or more of the predefined keywords such as "fake", "counterfeit", "knockoff", etc. The correlation matrix of Table 2 shows a consistency between the disclosing indicator and the fake indicator generated by topic modeling. A Likely Fake prediction can be discerned from both the absolute count and percentage of disclosing reviews. To this end, I define a dummy variable *Counterfeit_10_01*, which is set to 1 if a product has at least 10 disclosing reviews that account for more than 1% of the total number of reviews; it is set to 0 otherwise.

I build six of the most commonly used machine learning classifiers (such as naïve Bayes and random forest) with multiple groups of predictors according to their correlations to the labeled counterfeit dummy. Among all the review-generated topics, overall positive and negative sentiments (*Like, Dislike*), negative opinions on product authenticity (*Fake*), negative opinions on the fragrance's scent and durability (*Scent_negative*, *Last_negative*) are most informative in prediction. Features of rating score distribution (i.e., *1-star ratings percentage, 3-star ratings percentage and 5-star ratings percentage*) are also highly correlated with the propensity of an ASIN being labeled Likely Counterfeit. I optimize the classifier by adjusting variables and the average accuracy is raised up to 83% in the second model displayed in Table 4. I select the random forest classifier with the highest accuracy to generate the variable of interest, which is counterfeit likelihood, for all the products in our sample.

Table 4. Classification Model Accuracy

Model Predictors	(1) LR	(2) NB	(3) CART	(4) RF	(5) SVM	(6) LDA
Model I	0.71	0.68	0.72	0.79	0.70	0.72
Model II	0.70	0.73	0.71	0.83	0.69	0.73

Notes: LR = Logistic Regression, NB = Naïve Bayes, CART = Classification and Regression Trees, RF = Random Forest, SVM = Support Vector Machine, LDA = Linear Discriminant Analysis. The first column shows different combinations of predictors. Accuracy is defined as (# True Positives + # True Negatives) / # Cases. The variables in the two models are as follows. Model I has the predictors Overall, Disclose Ratio, Counterfeit_5_01, Vine, Fake_h, Dislike_h, Scent_positive, Last_positive, Price_positive, Review_1_star_pct, Review_3_star_pct, Review_4_star_pct. Model II is similar, except that Overall, Disclose Ratio, Counterfeit_5_01 and Last_positive are replaced by Overall_h, Num_Disclose, Counterfeit_10_01, and Last_positive_h, respectively.

2.3 Identification II – A Direct Measure Based on an Exogenous Event

2.3.1 The Main Direct Measure – Percentage Drop in the Number of Optional Sellers

Beyond developing the machine learning classification approach to capture perceived counterfeit probability from user-generated reviews, I provide a second identification method in this section, directly measuring counterfeit intensity based on an exogenous event. As introduced in Section 2.1, Amazon launched Project Zero to fight against online counterfeiting activities. This AI-based service consists of three parts: i. automated protections that scan the daily market to detect and ban suspicious items; ii. a self-service removal tool that allows brands themselves directly remove illegal copies; iii. An optional serialization service that assigns unique codes to items so that brands can apply them in the manufacturing or packaging process (Mehta, 2020). As more and more brands enroll in Project Zero, the first two services will gradually remove a lot of likely counterfeit products.

For those ASINs with multiple *optional offers (i.e., optional sellers)*, the effect will be manifested as a reduction in the number of optional offers. Naturally, the more intense counterfeit activities are under a specific listing, the greater drop in the number of optional sellers can be observed on average. Therefore, I use the standardized percentage drop in the number of optional sellers after the full adoption of Project Zero as the main direct measure of listing-level counterfeit intensity, defined as *D1*. Combining original brands' insights and knowledge of their own products and the platform's advanced deep learning technology, Project Zero provides the most reliable detection of online counterfeit products. Leveraging this event as an exogenous shock to the market also helps resolve the concern of lacking in ground truth and provides consistency in this study.

Specifically, I define three periods to better capture the effect of this anti-counterfeiting event. *Period I* starts from January 2018 and ends at December 2018, representing the phase when Project Zero was not launched and counterfeiting activities are not proactively regulated. Period II covers from January 2019 to December 2019, during which Project Zero was ready to be launched and gradually adopted by more and more brands on the US market. Period III covers from January 2020 to March 2020, referring to the phase when counterfeiting activities are initially mitigated after Project Zero and before the market was largely affected by the pandemic. The percentage drop of the average number of optional sellers from Period I to Period III is used to measure the counterfeit probability of each listing before Project Zero.

2.3.2 Data and Variables

In order to collect the optional sellers and sales information before 2019, I scraped historical Amazon data from Keepa, an Amazon research and pricing tracking platform. I focus on the men's fragrance category and search the 1037 top ranked listings specified in the previous identification method from the database. Similarly, I collect every historical update of their prices with shipping cost, sales ranks in the main category, and the number of optional offers, convert them to daily records, and aggregate them to monthly panels.

First, I calculate the average number of optional sellers in Period I (January 2018 to December 2018) and that in Period III (January 2020 to March 2020), and obtain the percentage drop in it from Period I to Period III. Specifically, the percentage drop is defined as follows:

$$Pct\_Drop = \begin{cases} \dfrac{No.offers\ in\ Period\ I - No.offers\ in\ Period\ III}{No.offers\ in\ Period\ I}, \\ if\ No.offers\ in\ Period\ I > No.offers\ in\ Period\ III \\ 0, \quad otherwise \end{cases} \tag{1}$$

If the percentage drop of one listing is less than zero, namely the number of offers in Period III is greater than that in Period I, it will be standardized as 0. Therefore, the standardized percentage drop in the number of offers is distributed between 0 and 1. The standardized percentage drop of the number of offers is used as the direct measure of counterfeiting intensity and will be included in the econometric model, denoted as *D1*. Second, to keep consistency with Identification I, I extract the same review metric variables based on the historical reviews up to the end of 2018, such as the word count, the image count, and the number of helpful votes. I generate aggregated average measures up to the current month within each listing to form the monthly panel. Last, the sales and prices data during Period I will be used to conduct the econometric analysis, combined with time-varying review metrics and time-invariant product features. Summary statistics of sales and prices data during Period I is presented in Table 5.

Table 5. Summary Statistics of Keepa Sales Data

VARIABLES	Description	(1) N	(2) mean	(3) sd	(4) min	(5) max
Share	Market Share (%)	8,219	0.133	0.364	0.00138	11.8
Sales	Estimated Sales of Item (count)	8,219	471.1	1,447	5.062	69,026
Sales Rank	Sales Rank in Main Category	8,219	150,846	196,863	81.21	3.224e+06
Price	Unit Price of Item ($)	8,219	40.00	30.87	3.980	284.7
Num_Reviews_v	No. User Reviews	8,219	132.7	271.0	1	2,653
Num_Helpful_Votes_v	Average No. Helpful Votes (per Review)	8,219	1.247	1.787	0	50
Num_Images_v	Average No. Images (per Review)	8,219	0.0193	0.0849	0	1.667
Text_Wordcount_v	Average No. Words (per Review)	8,219	20.92	11.68	1	220
Pct_Drop	Percentage Drop in the No. Offers	770	0.18831	0.39122	0	1
Disclose_Ratio_18	Percentage of Disclosing Reviews	770	0.01136	0.01705	0	0.13239

2.3.3 Model-Free Evidence

To confirm the impact of Project Zero on the number of optional offers within listings, a simple t-test is conducted to compare the average number of optional offers in 2018 (Period I) and that in the first three months of 2020 (Period III). The average number of optional offers among all the ASINs in our sample is 16.5 during Period I and 12.5 during Period III; the difference in between is statistically significant at the level of 0.001. Figure 4 shows how the average number of offers changed across the three periods, especially how it was gradually decreasing during Period II when Project Zero was launched and gradually took effects. After Project Zero, about 63% of the listings have seen some decrease in the number of offers; 19% of the listings lost at least half of their offers.

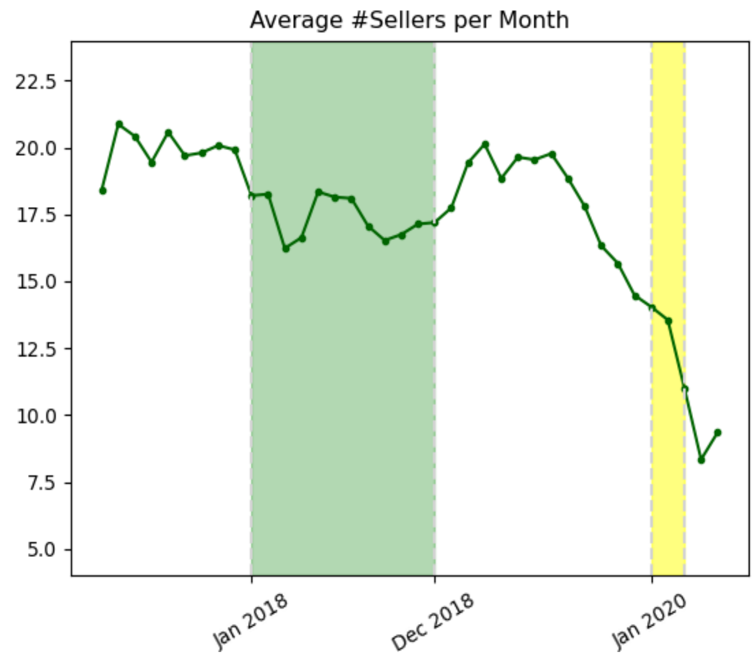

Figure 4. Average Number of Optional Offers

In Figure 5, the average number of optional offers is plotted separately for Likely Counterfeit listings and Likely Authentic listings specified in Identification I. Before and during Period I, there is a significant gap between the two groups of listings. Likely Counterfeit listings tend to have more optional sellers than Likely Authentic sellers, which coincides with our assumption and observation that counterfeit sellers tend to join crowded listings. After the intervention of Project Zero, the average number of optional offers of Likely Counterfeit listings dropped more greatly than that of Likely Authentic listings and the gap is narrowed in Period III. The significant difference of the two curves provides some consistency between Identification II and Identification I.

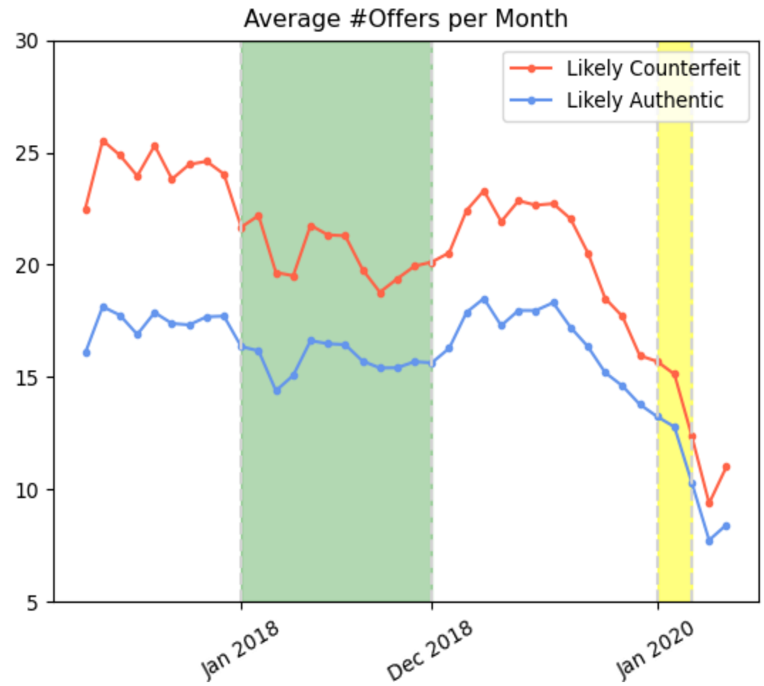

Figure 5. Average Number of Optional Offers – Likely Counterfeit vs. Likely Authentic

2.3.4 The Alternative Direct Measure – Percentage of Disclosing Reviews

To check the robustness of the direct measure, an alternative proxy of counterfeiting intensity based on user-generated contents is provided. As mentioned in Section 2.2, reviews including knockoff-related keywords are defined as *disclosing reviews*. I use the percentage of disclosing reviews among all historic reviews one listing has up to the end of Period I (*disclose_ratio_18*) as the second direct measure of counterfeiting activities, denoted as *D2*. As we have known from the Identification I, disclosing reviews are often diluted by many positive reviews. *Disclose_ratio_18* is distributed between 0 and 13.24% and half of the listings have at least one disclosing review.

In Figure 6, I plot the average number of offers for listings with and without disclosing reviews separately. For listings that have disclosing reviews, the average number of offers decreased from about 22 to 15; while for listings that haven't received any disclosing reviews, the average number of offers only dropped slightly from 11 to 10.

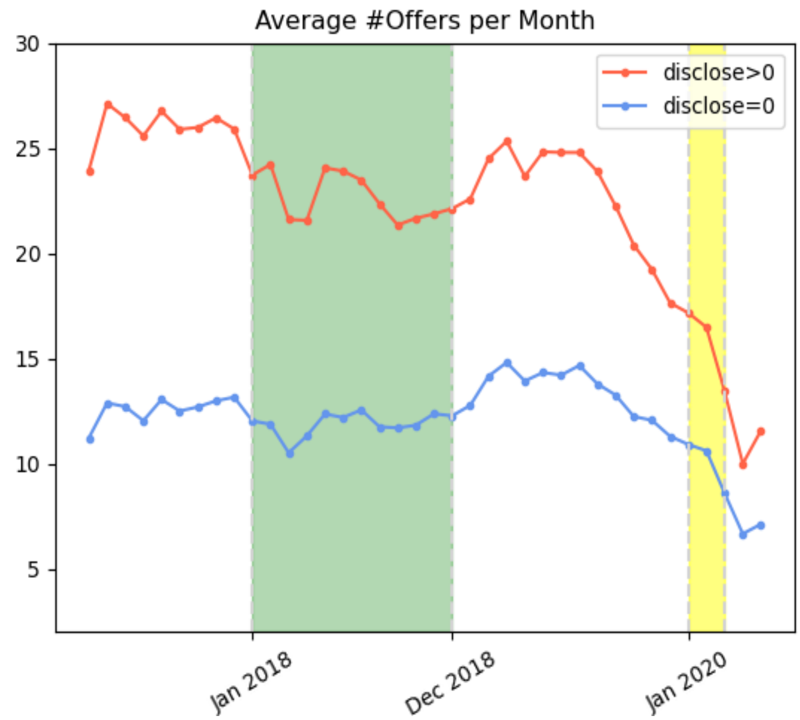

Figure 6. Average Number of Optional Offers – by Disclose Ratio

CHAPTER 3

Econometric Analysis on Consumer Choices and Market Shares

3.1 Discrete BLP Choice Model

3.1.1 Model Specification

I adopt a structural model of discrete choice with random coefficients, following Berry et al. (1995), to estimate the effect of perceived product authenticity on consumer choice and study what impact likely counterfeit products have on likely authentic ASIN sales, consumer utility and platform welfare. The BLP (Berry et al. 1995) model is a logit model estimating demand in differentiated product markets using aggregate market share data and allows for random coefficients of product characteristics and endogenous prices. I specify random coefficients for price and counterfeit probability, allowing consumer heterogeneity along these dimensions. I also expand the model by allowing for endogeneity of both price and counterfeit probability — the traditional BLP model only has endogenous price.

Specifically, the utility of consumer *i* buying a product in ASIN *j* in market *t* is defined as follows (we will use ASIN and product interchangeably):

$$u_{ijt} = \alpha_i P_{jt} + \gamma_i \, C_j + X_{jt}^{v} \beta^{v} + X_{j}^{inv} \beta^{inv} + \xi_{jt} + \varepsilon_{ijt}, \tag{2}$$

where *i* represents the consumer, *j* indexes the product, and *t* represents an Amazon fragrance market *t* (week *t* in our setting). P_{jt} is the weekly-average price (adjusted by discounts) of product *j* in market *t*, and C_j is the probability of *j* being a *likely counterfeit* product, generated from the machine learning classification model.[2] It also represents the consumer's perceived skepticism of the product's authenticity (buying from one of the sellers listed under the ASIN) after reading historical reviews online. X^v refers to time-varying product features such as rating valence and volume, numeric metrics extracted from the reviews such as average numbers of helpful votes and images. X^{inv} represents time-invariant product characteristics, such as the size and parfum concentration level in the case of fragrances. ξ_{jt} is the market-specific unobserved product attribute and ε_{ijt} is the random error, assumed to be i.i.d. type I extreme value distribution. The daily-level sales data is aggregated into weekly averages to construct a ten-week panel, which corresponds to the ten markets in our empirical analysis.

Although consumers are assumed to have vertical preferences on product authenticity, they are not equally familiar or engaged with the product review system, which creates heterogeneity in their sensitivity to signs of counterfeiting embedded in user-generated reviews. Consumers are also assumed to be heterogeneous in their price preferences, leading to the following specification for consumer distribution:

$$\binom{\alpha_i}{\gamma_i} = \binom{\bar{\alpha}}{\bar{\gamma}} + \Sigma v_i\,, \qquad v_i \sim N(0, I), \tag{3}$$

where v_i is consumers' unobserved preference for price and counterfeiting probability; in particular:

[2] Yang et al. (2022) discusses potential bias resulting from the correlation in measurement error of the predicted covariate (C_j) and the regression error. However, the remedy suggested by them is not applicable here, as we do not have access to the ground truth, as they do.

$$\alpha_i = \bar{\alpha} + \alpha_v v_i, \gamma_i = \bar{\gamma} + \gamma_v v_i. \tag{4}$$

This model can be viewed as a special case of the traditional BLP setup, wherein one of the observable product characteristics in which consumers have heterogeneous tastes — the counterfeit probability C_j — is time-invariant. Further, I allow C_j to be a second endogenous variable, in addition to price, as I further explain later. As indicated by Baum et al. (2002), the GMM estimator is still efficient when constructing instrumental variables for multiple endogenous variables, which allows us to follow the basic estimation process of the BLP paper. Accordingly, the probability that consumer *i* would choose product *j* is given by:

$$\begin{aligned} P\big(y_i = j | v_i, \varepsilon_{ijt}\big) \\ &= P\big(\alpha_i P_{jt} + \gamma_i C_j + X_{jt}^{v}\beta^{v} + X_j^{inv}\beta^{inv} + \xi_{jt} + \varepsilon_{ijt} \\ &> \alpha_i P_{lt} + \gamma_i C_l + X_{lt}^{v}\beta^{v} + X_l^{inv}\beta^{inv} + \xi_{lt} + \varepsilon_{ilt}, \forall l \neq j\big). \end{aligned} \tag{5}$$

Market share of product *j* is obtained by integrating over the i^{th} argument of the joint cumulative distribution function for each product $l \neq j$ (McFadden 1973), with the utility for the outside option normalized at 0. Following Berry et al. (1995), with the utility function containing two components of consumer heterogeneity, i.e., ε_{ijt} and v_i, the market share function can be obtained in two stages. First, integrating out over the ε_{ijt} conditional on v_i gives us a logit model (as in Equation 6 below), following McFadden (1973). Second, integrating out over v_i gives us the market shares as a function of product attributes. This second integration does not have a closed form, so Monte Carlo simulation agent data is used as a substitute in the estimation process (Berry et al. 1995).

$$s_{ijt} = Pr(y_{it} = j) = \frac{\exp\left(X_{jt}^{v}\beta^{v} + \alpha_i P_{jt} + \gamma_i C_j + X_j^{inv}\beta^{inv} + \xi_{jt}\right)}{1 + \sum_{l=1}^{J} \exp\left(X_{lt}^{v}\beta^{v} + \alpha_i P_{lt} + \gamma_i C_l + X_l^{inv}\beta^{inv} + \xi_{lt}\right)}. \quad (6)$$

3.1.2 Market Share and Instrument Variables

To obtain sales data to calculate market shares, I follow the approach widely used in previous research to convert sales rank into a proxy of sales (Chevalier and Goolsbee 2003, Ghose and Sundararajan 2006). To this end, product sales rank is assumed to follow a Pareto distribution: $Pr\ (s > S)\ = (k/S)^{\theta}$. For a particular product, the probability of randomly drawing a more popular competing item is taken to be equal to the number of items that are ranked ahead of the given product, which can be modeled as $(Rank - 1)/(Total\ number\ of\ items) = (k/S)^{\theta}$. Taking logs of the two sides transforms the equation into $ln\ (Rank - 1)\ = c - \ \theta * ln\ (Sales)$. Therefore, to convert rank data into actual sales, we only need to estimate the above log-linear regression model and obtain the coefficients c and θ, for which I conducted a simple experiment. First, two products are selected whose initial sales are ranked low enough to be approximately 0. Then I purchased a few copies of these products and observed how sales rank changed. I repeated the purchase and tracked the updated sales ranks several times within two days, collected data point pairs, and fit the aforementioned log-linear model. The estimated coefficient of men's fragrances product category is about 1.25, which falls between the suggested range of 0.9 to 1.3 in prior literature. In this way, we are able to estimate product weekly average market shares by feeding sales rank data into the model.

The traditional BLP model allows for endogenous prices and uses sums over product characteristics within or across brands as instrumental variables. Here, I select two characteristics of the men's fragrances category, size in ounces and volume (as a common measure of the concentration of alcohol and parfum in fragrances) and construct sums over characteristics of both non-rival goods (other goods under the same brand) and rival goods (goods of other brands) as instrumental variables for prices. In particular, I encode the volume from 0-4 according to the percentage of alcohol and parfum (i.e., after shave as 0, cologne as 1, eau de toilette as 2, eau de parfum as 3 and parfum as 4).

It should be noted that the distribution of fake products across ASINs is not random. Counterfeiters tend to target ASINs which are more expensive or popular, not only because they can make higher profits for higher cost products, but also because they have a greater chance to attract more consumers via a slight price reduction relative to authentic products. The figures shown in Section 2.2 also support the above intuition. This selection issue will bias the estimation of counterfeiting impacts. Therefore, I extend the traditional BLP model by allowing for endogeneity of both price and counterfeit probability.

The instrumental variables for counterfeit probability are generated as follows. I hypothesize that the likelihood of entry of fake sellers in an ASIN is increasing in the number of alternative sellers in that ASIN. This is due to the fact that hiding in a popular listing with many seller options reduces the risk of being reported by consumers, and consequently removed from the platform. Accordingly, I collect the number of multiple sellers (or buying options) along with the variance of listed prices as the first group of instruments for the counterfeit probability variable. Another group of instruments for counterfeit probability is derived from topic-modeling variables such as the positive or negative attitude on fragrances' scent and lasting power. These variables are likely to be correlated with the intensity of counterfeiting at the ASIN level, and relatively independent of the aforementioned selection issues. Table 6 shows the correlations between instruments and endogenous variables.

Table 6. IV-Relevance: First Step of 2SLS

VARIABLES	(1) Price	(2) Price	(3) Counterfeit Prob	(4) Counterfeit Prob
Size_other	-0.318045***	-0.186214***		0.001241
	(0.08971)	(0.067510)		(0.000982)
Size_rival	-0.0229577	-0.028248		-0.000509***
	(0.025745)	(0.020048)		(0.000157)
Vol1_other	0.323014	0.202809		-0.001183
	(0.491910)	(0.406074)		(0.004706)
Vol1_rival	-0.050305	-0.013544		0.003423***
	(0.273694)	(0.237968)		(0.001273)
Vol2_other	0.871385***	0.511794**		-0.004794
	(0.306181)	(0.243934)		(0.003386)
Vol2_rival	0.155685**	0.138580**		0.000990***
	(0.076206)	(0.067302)		(0.000374)
Vol3_other	5.545864***	3.166299***		0.016006***
	(0.597524)	(0.483957)		(0.005093)
Vol3_rival	-0.311296	-0.023869		0.003525***
	(0.235273)	(0.181020)		(0.001077)
Vol4_other	5.790964*	1.556321		0.032382
	(3.342363)	(3.026586)		(0.024023)
Vol4_rival	-2.562928	-2.328758		0.006293
	(2.004261)	(1.915891)		(0.005166)
Num_Options		-0.617933***	0.001594**	0.001840***
		(0.055281)	(0.000642)	(0.000638)
Price S.D.		2.865231***	0.005654***	0.003621***
		(0.227237)	(0.001279)	(0.001140)
Scent_positive		-17.83402	-0.594439***	-0.561581***
		(14.63676)	(0.122517)	(0.123626)
Last_positive		31.74136***	-0.522849***	-0.502656***
		(11.81995)	(0.092974)	(0.092838)
Scent_negative		-17.5099	-0.016489	0.035845
		(18.0219)	(0.152118)	(0.154955)
Last_negative		-0.255102	-0.009920	-0.036801
		(5.855067)	(0.052059)	(0.051983)
Price_positive		-17.08565	-0.037505	-0.017014
		(13.93445)	(0.113698)	(0.113795)
Constant	77.01404***	61.14638***	0.478586***	0.497944***
	(17.09209)	(19.27038)	(0.035243)	(0.056164)
Observations	9,520	9,520	9,520	9,520
R-squared	0.135	0.432	0.082	0.122
F-statistic	27.01	41.88	14.78	9.62

Notes: Columns (1) – (2) test the correlation between IVs and the first endogenous variable — Price; Columns (3)-(4) test the correlation between IVs and the second endogenous variable — Counterfeit Probability. Standard errors in parentheses. *** $p<0.01$, ** $p<0.05$, * $p<0.1$.

3.2 Empirical Results with Identification I

3.2.1 Main Results

I first estimate a reduced-form 2SLS fixed-effects model to explore the effects of prices and counterfeit probability on market shares and the efficacy of instrument variables. Next, following the implementation of Vincent (2015), I include consumer heterogeneity and estimate the BLP random coefficient logit model (Equations 2-5) built on simulated data generated from Monte Carlo analysis. Last, I run a sub-sample analysis on a data set containing only top-ranked products whose market shares are greater than 0.1% to examine if there is a different pattern associated with the most popular products.

Table 7 shows the results of the 2SLS models; Columns 1-3 correspond to the top 1037 products whose market shares are no less than 0.01%, with 9520 listing-week observations. The regressions include brand fixed effects and week fixed effects. Columns 4-6 report estimations of the same models on the top 148 products whose market shares are no less than 0.1%, resulting in 807 listing-week observations. The coefficient estimates for the primary variables are consistent across the models. Demand is negatively correlated with price. The coefficient of counterfeit probability estimated on the larger data set is negative and significant, as expected. We also find a significant positive coefficient on the log of number of ratings, suggesting a positive effect of popularity on consumers' purchasing decisions. The coefficients of *Image_count* and *Helpful_votes* are also significantly positive, indicating that the average product quality suggested by historical reviews is positively associated with consumer utility and product sales.

Table 7. 2SLS Estimation Results: Second Step of 2SLS

VARIABLES	(1)	(2)	(3)	(4)	(5)	(6)
Price	-0.000256*	-0.000361**	-0.000261*	-0.001186	-0.002350	-0.001379
	(0.000146)	(0.000179)	(0.000151)	(0.001693)	(0.001866)	(0.001895)
Counterfeit Prob	-0.106375***	-0.110953***	-0.120574***	-0.709483***	-0.920389***	-0.622347***
	(0.0364604)	(0.042532)	(0.039091)	(0.194334)	(0.216412)	(0.195610)
Log Num_Reviews	0.047042***	0.047410***	0.047906***	0.416064***	0.418394***	0.413630***
	(0.003688)	(0.004125)	(0.003866)	(0.035293)	(0.035976)	(0.034794)
Rating	-0.022794*	-0.021920	-0.024357*	0.413180**	0.34367	0.492899***
	(0.013817)	(0.014861)	(0.014146)	(0.187580)	(0.213166)	(0.189710)
Num_Images	0.086860***	0.093487***	0.085741***	2.473887***	1.824246**	2.729471***
	(0.032577)	(0.033524)	(0.032667)	(0.847526)	(0.888416)	(0.830697)
Num_Helpful_Votes	0.004315***	0.004489***	0.004466***	-0.007708	-0.002683	-0.006420
	(0.001564)	(0.001573)	(0.001567)	(0.006333)	(0.006642)	(0.006205)
Text Wordcount	0.000107	0.000104	0.000101	0.009321	0.005062	0.011025
	(0.000308)	(0.000310)	(0.000310)	(0.011315)	(0.011588)	(0.011103)
Shipping			-0.007122			-0.305890***
			(0.005522)			(0.076807)
Size		0.001840	0.001703		0.028323	0.010703
		(0.001478)	(0.001464)		(0.018709)	(0.017868)
Vol 1		0.037564				
		(0.037653)				
Vol 2		0.035949			0.286329***	
		(0.037548)			(0.096964)	
Vol 3		0.045307			0.176236	
		(0.038699)			(0.135462)	
Vol 4		0.036580				
		(0.048751)				
Constant	-0.077656	-0.12235	-0.056983	-4.251418***	-4.184785***	-3.941059***
	(0.062728)	(0.076561)	(0.064024)	(0.951118)	(1.057384)	(0.956755)
Observations	9,520	9,520	9,520	807	807	807
R-squared	0.034	0.032	0.031	0.155	0.131	0.189

Notes: Columns (1) – (3) are based on the complete data set, while Columns (4) – (6) are based on the subset of top ranked listings whose market shares are no less than 0.1%. The dependent variable is market share. Time period is Feb-Apr 2021. All regressions include time (week) fixed effects and tier (seven brand categories) fixed effects. Standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

We now discuss the results from our main random coefficients (BLP) logit model, shown in Table 8a. The dependent variable is the market share and the latent variable is consumer mean utility. Columns 1-2 correspond to the full data set, whereas Columns 3-4 are for products whose market shares are no less than 0.1% (approximately the Top 20th percentile in average sales). We see that a higher counterfeit probability significantly reduces consumer mean utility, and consequently reduces sales and market share as well. The effect of price on consumer mean utility is also negative. On the other hand, the coefficients of rating and log of rating counts are both significantly positive, suggesting a better reputation and sales history have positive impacts on consumer utility. Turning to Columns 3-4, comparing these results with those for the larger data set, we see that the effects of prices and some review metrics on mean utility are no longer significant, while the magnitude of counterfeit probability is higher. This is consistent with the notion that consumers who seek to consume from best sellers attach relatively more importance to authenticity and historical sales record when they make purchasing decisions.

Table 8b shows the estimated standard deviations of random coefficients based on consumers' heterogeneous preferences, with the four columns corresponding to the models presented in Table 8a, respectively. Looking at Column 2 (the complete model), we see that the mean effect of counterfeit probability is -1.0073 and standard deviation is 2.6599, indicating that about 64.90% of the mass of the distribution is in the negative range (assuming a normal distribution). The estimates are more accurate for best-selling products (Column 4), and the proportion of the distribution in the negative range is 89.04%.

Table 8a. Results for BLP Choice Model with Random Coefficients – Identification I

Variables	(1)	(2)	(3)	(4)
Price	-0.003699***	-0.006356***	-0.000610	-0.001654
	(0.000794)	(0.001028)	(0.001783)	(0.001878)
Counterfeit Prob	-1.030453*	-1.017308*	-2.277378**	-2.315536**
	(0.593504)	(0.611455)	(1.158386)	(1.04508)
Rating	0.872719***	0.930219***	-0.387738**	-0.709886***
	(0.055122)	(0.060511)	(0.191929)	(0.213216)
Log Num_Reviews	0.578421***	0.586222***	0.645247***	0.643096***
	(0.014960)	(0.015585)	(0.039647)	(0.038746)
Num_Images	1.491585***	1.827954***	1.452147*	-0.505768
	(0.129598)	(0.138797)	(0.775719)	(0.911570)
Num_Helpful_Votes	0.058544***	0.068111***	0.025080***	-0.014852
	(0.005441)	(0.005726)	(0.006142)	(0.011865)
Text Wordcount		-0.006278***		0.025713***
		(0.001168)		(0.006658)
Size		0.017985***		0.037167**
		(0.005377)		(0.018669)
Vol 1		0.281208**		
		(0.137081)		
Vol 2		0.022369		-0.008987
		(0.137109)		(0.097472)
Vol 3		0.460014***		-0.379501***
		(0.142057)		(0.135343)
Vol 4		0.411517**		
		(0.180481)		
Brand Category FE	Yes	Yes	Yes	Yes
Observations	9520	9520	807	807

Notes: Columns (1) – (2) are based on the complete data set, while Columns (3)–(4) are based on the subset of top- ranked listings whose market shares are no less than 0.1%. The dependent variable is market share and the latent variable is consumer mean utility. Time period is Feb-Apr 2021. All regressions include tier (brand category) fixed effects. Vol 1, Vol 2, Vol 3, Vol 4 are dummy variables that indicate whether a perfume is Eau de Cologne, Eau de Toilette, Eau de Parfum or Parfum, respectively. Robust standard errors in parentheses. *** $p<0.01$, ** $p<0.05$, * $p<0.1$.

Table 8b. Standard Deviation of Individual Random Coefficients – Identification I

Variables	(1)	(2)	(3)	(4)
SD — Price	0.000066	0.000093	2.70e-07	7.53e-11
	(0.065769)	(0.059590)	(0.073156)	(0.069797)
SD — Counterfeit Prob	2.644236***	2.659966***	2.205202**	1.88428**
	(0.599278)	(0.601004)	(1.003856)	(0.940913)

Notes: Columns (1) – (4) correspond to those in Table 8a. Robust standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

3.2.2 Price Elasticities

We now study economic impacts of counterfeiting on merchant profits and platform welfare, by generating own and cross-price elasticities, for likely counterfeit and likely authentic products, based on the BLP model results. Recall that products in our sample are labeled as *likely counterfeit* or *likely authentic* based on the counterfeit probabilities generated by the classification model, using a 50% cutoff. I examine the elasticities across the ten markets (weeks) in the data set and also across seven horizontal fragrance brand categories, as follows: Low-End Fashion brand (e.g., Abercrombie & Fitch); High-End Fashion brand (e.g., Davidoff); Designer's Fragrances brand (e.g., Paco Rabanne); Luxury brand (e.g., Chanel); beauty brand (e.g., Lancome); Price-Friendly brand (e.g., Bod Man); and Auto brand (e.g., Mustang). Among the seven categories, High-End Fashion (2), Designer's Fragrances (3) and Luxury (4) brands sell perfumes at relatively higher prices. Table 9 shows own and cross-price elasticities for the ten markets, while Table 10 shows the same across the brand categories.

Starting with own-price elasticities in Table 9, the average value of own price-elasticity of likely counterfeit products is -0.3186 (i.e., a 1% increase in the price of a likely counterfeit product is associated with a 0.3186% decrease in its own market share on average. The average own price-elasticity of likely authentic products is a bit lower, at -0.2421. It is natural that there is more inelastic demand for likely authentic products, as users are relatively less price sensitive to products that are likely to be authentic as compared to listings that are likely to yield a counterfeit purchase. Looking at own-price elasticities across brand categories, we once again see that demand is relatively more inelastic for likely authentic products. Across brand categories demand is most elastic for the expensive Beauty and Luxury brands and most inelastic for the cheaper Price-Friendly and Low-End Fashion categories.

Table 9. Own and Cross-Price Elasticities Across Ten Markets

	(1) Own Price Elasticity Counterfeit	(2) Own Price Elasticity Authentic	(3) Cross-Price Elasticity Counterfeit on Authentic
Market 1	-0.323968***	-0.241899***	0.000113***
	(0.013164)	(0.000299)	(0.000002)
Market 2	-0.320058***	-0.242565***	0.000117***
	(0.013451)	(0.000321)	(0.000000)
Market 3	-0.320543***	-0.241638***	0.000107***
	(0.012869)	(0.000297)	(0.000001)
Market 4	-0.318547***	-0.239869***	0.000109***
	(0.011597)	(0.000278)	(0.000001)
Market 5	-0.317793***	-0.239948**	0.000103***
	(0.011597)	(0.000276)	(0.000001)
Market 6	-0.309141***	-0.243010***	0.000116***
	(0.012059)	(0.000294)	(0.000001)
Market 7	-0.314103***	-0.240456***	0.000105***
	(0.012030)	(0.000299)	(0.000001)
Market 8	-0.317099***	-0.241858***	0.000103***
	(0.011472)	(0.000299)	(0.000001)
Market 9	-0.319383***	-0.243340***	0.000109***
	(0.011727)	(0.000307)	(0.000001)
Market 10	-0.325582***	-0.246482***	0.000128***
	(0.013446)	(0.000416)	(0.000002)

Notes: The elasticities are estimated in the BLP random coefficient model. Robust standard errors are in parentheses. *** p<0.01, ** p<0.05, * p<0.1

Turning to cross-price elasticities, in Table 8, the average (across the 10 markets) cross-price elasticity between likely counterfeit and likely authentic products is 0.00011, so that for a randomly chosen pair of likely counterfeit and likely authentic products, a 10% decrease in the price of the likely counterfeit product is associated with a 0.0011% decrease in the market share of the likely authentic product, on average. While the magnitude of the average cross-price elasticity appears to be small, it should be noted that this is a diverse market with thousands of products, so that the cumulative effect of systematic price reductions by counterfeit sellers on the market shares of authentic sellers could be substantive in the aggregate. This issue will be further examined in the counterfactual experiments of Chapter 4.

Looking at cross-price elasticities between likely authentic and likely counterfeit products across brand categories, in Table 10 and Figure 7, the average cross price-elasticity between likely counterfeit and likely authentic products is the lowest in the Price-Friendly brand category, which is 0.000018. And the Low-End Fashion (1) category has the highest cross price-elasticity between likely counterfeit and likely authentic products, which is 0.000293 (i.e., a 1% decrease in the price of a likely counterfeit product in a Low-End Fashion brand is associated with a 0.000293% decrease in the market share of a likely authentic product of the same brand category). Besides, the average cross price-elasticities between likely counterfeit and likely authentic products within High-End Fashion, Designer's Fragrances and Luxury brands are also relatively high.

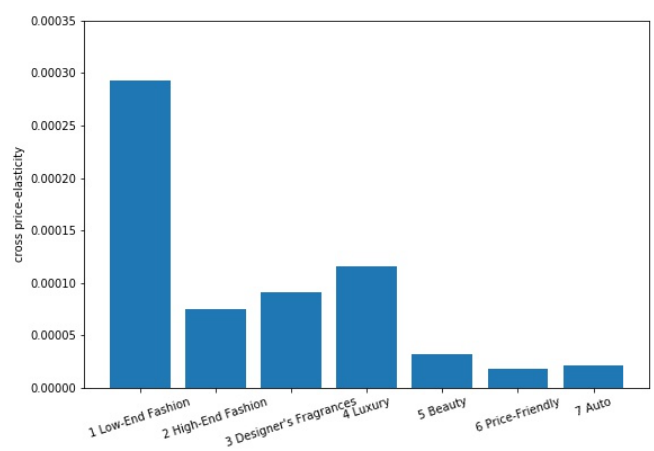

Figure 7. Average Cross-Price Elasticities Across Brand Categories

Table 10. Own and Cross-Price Elasticities Across Brand Categories

	(1) Own Price Elasticity Counterfeit	(2) Own Price Elasticity Authentic	(3) Cross-Price Elasticity Counterfeit on Authentic
Fashion_low	-0.223566***	-0.159890***	0.000293***
	(0.009829)	(0.002876)	(0.000007)
Fashion_high	-0.280193***	-0.262640***	0.000075***
	(0.007227)	(0.005673)	(0.000000)
Designer	-0.282122***	-0.268660***	0.000091***
	(0.005500)	(0.007165)	(0.000000)
Luxury	-0.404167***	-0.386510***	0.000116***
	(0.006469)	(0.005465)	(0.000000)
Beauty	-0.554014***	-0.227800***	0.000032***
	(0.035741)	(0.007822)	(0.000000)
Friendly	-0.151891***	-0.130790***	0.000018***
	(0.005192)	(0.001768)	(0.000000)
Auto	-0.187630***	-0.199780***	0.000021***
	(0.012209)	(0.007912)	(0.000000)

Notes: The elasticities are estimated in the BLP random coefficient model. Robust standard errors are in parentheses. *** $p<0.01$, ** $p<0.05$, * $p<0.1$

The positive price-elasticity between likely counterfeit and likely authentic products is consistent with a substitution effect, which appears to dominate any potential promotional effect for merchants of genuine products. This finding is qualitatively different from that documented for the traditional retail industry by Qian et al. (2014), in which advertising effects have been shown to dominate substitution effects for high-end products. The different outcomes have to do with the differences between non-deceptive versus deceptive counterfeit products. First, unlike the offline sales of luxury counterfeits which are targeted at a separate segment of consumers inclined to knowingly purchase cheap knockoffs, online deceptive knockoffs pretend to be authentic and are targeting the vast majority of consumers that are not looking for knockoffs. Second, unlike piracy of information goods or luxury copycats, prices of online knockoffs are not necessarily lower than that of authentic ones. Third, online consumers' perception of a counterfeit purchase happens only after the purchase is finished. Therefore, online counterfeiters take away larger market shares and profits from authentic manufacturers, and substitution effects dominate potential positive effects.

3.3 Empirical Results with Identification II

In this section, I discuss the results of econometric analysis when using Identification II – the direct measure – to capture the intensity of online counterfeiting activities. As aforementioned, there are two different direct measures: D1 (*Pct_Drop*), the percentage drop of the average number of optional offers, is based on an exogenous event and reflects the combination of original brands' insight and the platform's detection decision over authenticity; D2 (*Disclose_Ratio_18*), the percentage of disclosing reviews, is based on user-generated contents regarding the true quality they have perceived after the purchase. In this section, the predicted counterfeit probability C in Equation 2 will be replaced with D1 and D2 respectively, for the main analysis and the alternative analysis. In the main model, D1 will be used to measure the variable of interest – the likelihood a consumer encounters a fake purchase under a given ASIN, as it encloses the information and experts' judgement closest to the ground truth. Then, I will estimate the same BLP model with the measure replaced with D2 to provide complementary results, as it did show consistency in the identification of counterfeit listings despite the different scales. Note that the analyses with Identification II (both D1 and D2) are conducted on the data during January to December 2018, with twelve monthly periods, before the launch of Project Zero.

Similarly, consumers are assumed to browse the market and choose the listing that can maximize her utility. They are allowed to have heterogeneous preferences for the price and heterogenous level of capability to detect knockoffs. However, different from Identification I where the predicted counterfeit probability is considered endogenous, D1 is specified as an exogenous variable, since the launch of Project Zero is an exogenous event that does not correlate with product quality or other unobserved factors.

Table 11a shows the estimated linear coefficients of the random choice model and Table 11b shows the estimation of the random effects of stochastic variables. The dependent variable is the market share, and the latent variable is consumers' mean utility. Column 1 corresponds to the main model, using the exogenous D1 as the counterfeit measure. As we can see, the estimated effect of the price on consumers' mean utility is -0.007, suggesting one dollar's increase in the price will on average reduce consumers' utility by 0.00695, or the odds of the listing's market share by 0.7%. The estimated coefficient of *Pct_Drop* is -0.884, meaning every one extra percent (absolute) increase in the counterfeit measure will one average reduce consumers' mean utility by 0.00884. If the percentage drop of the number of offers converts from 0 to 1, namely an authentic listing becoming a counterfeit listing, the market share odds of that corresponding listing will significantly drop by 58.69%. Generally speaking, the counterfeit intensity will significantly hurt consumer's utility and therefore the listing's market share. Table 11b shows the standard deviation of the stochastic variables. As described in Section 3.1, consumer preferences for prices and counterfeit measures are assumed to be normally distributed. Combining Table 11a and 11b, we can conclude that 66.09% of consumers are negatively affected by a listing's counterfeit probability. This distribution of consumers' attitudes towards online counterfeit activities is quite consistent with what we have estimated in the main model of Identification I.

Column 2 shows results of the alternative model, using the exogenous D2 as the counterfeit measure. The coefficients of the price and the percentage of disclosing reviews are both significantly negative, indicating a consistent conclusion that the price and counterfeit activities will hurt consumers' utility on average and reduce the sale and market share of the corresponding listing. When combined with the results in Table 11b, we can derive the distribution of consumer attitudes towards counterfeiting signals embedded in user-generated reviews. About 77.28% of the individual coefficients of D2 are at the negative range, which is close to the random effect distribution of D1.

Table 11a. Results for BLP Choice Model with Random Coefficients – Identification II

Variables	(1)	(2)	(3)	(4)
Price	-0.006952***	-0.004493**	-0.021929***	-0.018113**
	(0.002070)	(0.001756)	(0.005350)	(0.008979)
Pct_Drop	-0.883521**		-0.544697*	
	(0.355630)		(0.330612)	
Disclose_Ratio_18		-26.89978**		-45.63434***
		(12.10556)		(16.33979)
Rating	0.606484***	0.628090***	1.146906***	0.855486***
	(0.071271)	(0.065076)	(0.145729)	(0.145809)
Log Num_Reviews_v	0.435641***	0.483972***	0.401049***	0.628101***
	(0.002963)	(0.021908)	(0.904983)	(0.032351)
Num_Images_v	0.330549***	0.329018***	0.726554***	0.219415*
	(0.069545)	(0.066760)	(0.082328)	(0.124312)
Num_Helpful_Votes_v	0.072942***	0.69696***	0.046575***	0.047393**
	(0.008012)	(0.008254)	(0.007585)	(0.020224)
Text_Wordcount_v	-0.012121***	-0.011895***	-0.010227***	-0.004958**
	(0.000593)	(0.000521)	(0.000669)	(0.002245)
Size	0.004163	-0.000468	-0.033573***	-0.047393***
	(0.004858)	(0.004179)	(0.006528)	(0.008837)
Vol 1	0.173064***	0.13066***	0.506287***	0.744166***
	(0.051042)	(0.039928)	(0.024583)	(0.075682)
Vol 2	0.035272	0.057418	0.253264***	0.649053***
	(0.043345)	(0.038264)	(0.028762)	(0.077991)
Brand Category FE	Yes	Yes	Yes	Yes
Observations	8219	8219	8219	8219

Notes: Columns (1) and (3) use D1 (Pct_Drop) as the variable of interest, while Columns (2) and (4) use D2 (Disclose_Ratio_18) as the variable of interest. In Columns (1) – (2), D1 and D2 are treated as exogenous; in Columns (3) – (4), D1 and D2 are treated as endogenous. The dependent variable is the market share and the latent variable is consumer mean utility. Time period is Jan-Dec 2018. All regressions include tier (brand category) fixed effects. Vol 1, Vol 2 are dummy variables that indicate whether a perfume is Eau de Cologne and Eau de Toilette, respectively. Robust standard errors in parentheses. *** $p<0.01$, ** $p<0.05$, * $p<0.1$.

Table 11b. Standard Deviation of Individual Random Coefficients – Identification II

Variables	(1)	(2)	(3)	(4)
SD — Price	0.008428***	0.007546***	0.026227***	0.029892***
	(0.002942)	(0.002511)	(0.003876)	(0.003170)
SD — Counterfeit Prob	2.131226***	35.9636***	4.238332***	0.000145
	(0.468949)	(10.0627)	(0.484639)	(1637.192)

Notes: Columns (1) – (4) correspond to those in Table 11a. Robust standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

I conduct robustness checks by converting direct measures D1 and D2 to endogenous variables. Following the design in Section 3.1, I constructed the monthly average number of optional offers and its square term during Period I as the instrument variables for endogenous D1 or D2. Estimation results are displayed in Table 11 Columns (3) – (4). Significantly negative effects of the price and the counterfeit measures are confirmed and consistent. Beyond that, we can also observe significantly positive effects of rating volume and valence, the average number of images attached in reviews, and the average helpful votes reviews have received across all or most models. On the contrary, the average length of review texts will negatively affect consumers' utility and listings' market shares, which could possibly be explained by its correlation with review frauds. In terms of product features, consumers tend to purchase eau de toilettes and colognes more often among multiple perfume types and prefer smaller packaging than large bottles.

3.4 Robustness Examination on a Utilitarian Good

To explore the impact of counterfeiting in another product category, and to validate the robustness of our methodology and findings, I conduct the entire analysis on a utilitarian product, that is a product designed to be objectively functional rather than subjectively attractive. Specifically, I look at cell-phone wireless chargers, a stereotypical utilitarian product, yet one that is subject to active counterfeiting. It is intrinsically different from perfumes in a number of respects. First, consumer preferences for wireless chargers are more homogeneous than that of fragrances, as there is inherently less demand for variety. More specifically, preferences are much more vertical rather than horizontal, in that consumers essentially care about quality rather than other subjective non-functional product characteristics. As a result, a counterfeit wireless charger tends to hurt consumer utility explicitly for not satisfactorily performing the charging function. Second, when it comes to wireless chargers, people care less about brand value and personal taste, as they do for high-end fragrances. Lastly, as consumers have a higher level of acceptance to chargers of less-established brands, we anticipate less counterfeit activity in the charger category as compared to fragrances. On Amazon, it is observed that about half of the wireless-charger related ASINs list a single seller; for those listings with multiple sellers, the average number of sellers is less than 5, which is substantially smaller than in the case of fragrances. Therefore, I only rely on Identification I to define the counterfeit intensity for cell phone wireless chargers. This also allows us to streamline the set of instruments by dropping the ones related to multiple sellers, as we explain below.

I define 15 counterfeit-related topics for cell-phone wireless chargers: counterfeit warning (fake), overall sentiments, attitudes towards price, shipping and quality in terms of charging speed, connection, flexibility, lifespan, design, etc. (positive or negative). Topic indicators are extracted using anchored correlation explanation topic modeling. Variables on authenticity, overall sentiment, charging speed, lifespan, compatibility, shipping, and services as well as rating distribution metrics are selected to train a random forest classifier. To construct the training data set, 200 (of the 565) listings in our data sample are manually labeled as likely counterfeit or likely authentic.

As mentioned earlier, when estimating the counterfeiting effect in cell-phone wireless chargers, one substantial difference from the prior analysis of fragrance product category is that most chargers have only one seller per ASIN, rather than multiple. Therefore, I exclude the number of seller options and standard deviation of price from the set of instrument variables. Instead, we include topic variables regarding authenticity, charging speed, and lifespan as instruments for counterfeiting probability. Compared to fragrances where one listing is often linked to over a dozen sellers, wireless chargers can be considered a special and simpler case where the number of options is one. Accordingly, the predicted counterfeit probability suggests the likelihood of a product (instead of a listing) being counterfeit. In addition, I define characteristics on multi-charging design (single charging, two-in-one, three-in-one, or four-in-one), and station design (pad, stand, or station) and generate sums of such characteristics over products within or across brands as the BLP-style instruments for endogenous prices. Lastly, I conducted purchasing experiments to convert charger sales rank into charger market share, as we did for fragrances.

The results for the charger category are reported in Tables 12a-b, which suggest that the counterfeit probability significantly hurts consumers' utility. A higher price reduces consumers' willingness to buy, while rating, images in reviews and free shipping have positive impacts on sales. Three-in-one and four-in-one charging stations are more popular than single or two-in-one chargers. While the results are quite consistent with what we found for men's fragrances, it is worth noting that the standard deviation of the individual price coefficient is significant in the case of cell-phone wireless chargers, indicating that user preferences for prices are more heterogeneous for the charger as compared to the fragrance category. Also, the individual coefficients of counterfeiting probability follow a normal distribution with estimated mean -2.788421 and estimated standard deviation 2.082199 (Column 4), suggesting that 91.97% of the distribution is in the negative range, which is substantially higher than the corresponding figure for fragrances.

Table 12a. Results for BLP Choice Models with Random Coefficients – Chargers

Variables	(1)	(2)	(3)	(4)
Price	-0.073162**	-0.057889**	-0.068125**	-0.044622***
	(0.030748)	(0.026254)	(0.028460)	(0.023549)
Counterfeit Prob	-2.402878***	-2.932925***	-2.330667***	-2.788421***
	(0.532091)	(0.532198)	(0.56752)	(0.572094)
Rating	1.459121***	1.265042***	1.437825***	1.251411***
	(0.193070)	(0.152019)	(0.189578)	(0.149357)
Log Num_Reviews	0.891418***	1.012251***	0.880677***	1.016091***
	(0.059620)	(0.023047)	(0.059236)	(0.022563)
Num_Images	1.89168***	1.027836***	1.812292***	0.972098***
	(0.289964)	(0.218348)	(0.289238)	(0.215449)
Text Wordcount	-0.010699***	-0.009009***	-0.010476***	-0.009594***
	(0.003462)	(0.002677)	(0.003442)	(0.002664)
Num_Helpful_votes	0.070377***	0.062240***	0.061425***	0.051032***
	(0.013909)	(0.013753)	(0.012917)	(0.0126382)
Log Num_Q&A	0.323801***		0.338438***	
	(0.065155)		(0.065379)	
Shipping	1.351283***	1.272033***	1.164859***	1.000808***
	(0.368550)	(0.327960)	(0.320866)	(0.278833)
3-in-1			0.252676**	0.253128**
			(0.134349)	(0.122817)
4-in-1			0.262263	0.273453*
			(0.163945)	(0.148172)
Pad			-0.247039**	-0.173943**
			(0.089228)	(0.080657)
Stand			-0.033535	-0.048918
			(0.091326)	(0.078858)
Brand Category FE	Yes	Yes	Yes	Yes
Observations	4462	5176	4462	5176

Notes: Columns (1) – (4) are based on the complete data set. The dependent variable is the market share and the latent variable is consumer mean utility. Time period is Dec 2020 to Apr 2021. 3-in-1 and 4-in-1 are dummy variables to indicate whether the product is a 3-in-1 or 4-in-1 charging station. Pad and Stand are also dummy variables referring to different styles of the charger. Robust standard errors in parentheses. *** $p<0.01$, ** $p<0.05$, * $p<0.1$.

Table 12b. Standard Deviation of Individual Random Coefficients – Chargers

Variables	(1)	(2)	(3)	(4)
SD — Price	0.037231***	0.028334**	0.036073***	0.023457
	(0.014499)	(0.014462)	(0.013971)	(0.015174)
SD — Counterfeit Prob	2.617766***	2.313474***	2.47865***	2.082199***
	(0.468727)	(0.387850)	(0.446156)	(0.372270)

Notes: Columns (1) – (4) correspond to the columns in Table 12a. Robust standard errors are in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

CHAPTER 4

Economic Impacts and Platform Strategies

4.1 Counterfactual Experiments – Design and Simulation

4.1.1 Experiment Design

To gain deeper insights into the impacts of counterfeiting on consumers, the market and platform, I conduct counterfactual experiments that leverage our structural parameter estimates from the random coefficient choice model. Specifically, three counterfactual scenarios or treatments are considered. In contrast to these treatments, the original data set used in previous sections is considered to be the *Control* group.

In Treatment I — *Counterfeit Shock* — we examine the impact of a proliferation of counterfeit products in the market. I simulate this scenario by shifting the distribution of counterfeit probability, in a stochastically dominant manner, so that the perceived probability of authentic product is uniformly reduced by 20%. In other words, the counterfeit probability is transformed from p, say, to $p + 0.2*(1-p)$. This will enable us to explore how the sales of likely authentic products as well as platform revenues are affected when counterfeit activity is ramped up on the platform.

In Treatment II — *Counterfeit Detection* — we consider the scenario where the platform applies detection algorithms of its own and signals the results in some form to consumers — maybe as a flag, banner, or other message in the ASIN listing. This is implemented by polarizing the distribution of counterfeit probability to increase the contrast between likely counterfeit and likely authentic products. Specifically, I increase the counterfeit probability by 50% if it is greater than 0.5 and decrease it by 50% if it is less than 0.5. This would serve to enhance the consumers' perception of counterfeit activity, affecting their choices and thereby market shares of products and platform revenues.

Finally, in Treatment III — *Counterfeit Ban* — we consider the scenario where the platform explicitly bans all likely counterfeit listings. That is, all listings with counterfeit probability of 0.5 and higher are eliminated from the marketplace. This case would reveal what would happen to consumer and platform welfare if the platform were to eliminate suspect listings, in order to weed out and deter counterfeiting activity.

4.1.2 Simulation

To implement the experiments and compare results of the treatment versus control groups, I simulate the decision making of 100,000 consumers, based on the estimation process of the BLP model, in three steps. First, we draw the random part of the price ($\alpha_v v_i$) and counterfeit probability coefficient ($\gamma_v v_i$), as in Equation (4), along with the unknown error ε_{ijt} in each market. ν_i are normally distributed by assumption; α_v and γ_v are the estimated standard deviations of the variables (Table 8b).[3] ε_{ijt} are i.i.d. and follow the Type I extreme value distribution, by assumption. Second, we compute the unobserved product characteristics ξ_{jt} (Berry 1995), as follows. Based on Equation (6), the market shares of product *j* and outside alternative 0 are (respectively):

[3] Note that for the category of men's fragrances, estimation of individual heterogeneity in prices (α_v) is not significant, so we only include the random effect of counterfeiting probability in the first step of simulation for this category.

$$S_{jt} = \frac{\exp\,(X_{jt}^{v}\,\beta^{v} + \bar{\alpha} P_{jt} + \bar{\gamma} C_j + X_j^{inv} \beta^{inv} + \xi_{jt})}{1 + \sum_{l=1}^{J} \exp\,(X_{lt}^{v}\,\beta^{v} + \bar{\alpha} P_{lt} + \bar{\gamma} C_l + X_l^{inv} \beta^{inv} + \xi_{lt})}, \tag{7}$$

$$S_{0t} = \frac{1}{1 + \sum_{l=1}^{J} \exp\,(X_{lt}^{v}\,\beta^{v} + \bar{\alpha} P_{lt} + \bar{\gamma} C_l + X_l^{inv} \beta^{inv} + \xi_{lt})}. \tag{8}$$

Using these we can solve for the unobserved product characteristics (see, e.g., Ghose et al. 2012):

$$\xi_{jt} = log\big(S_{jt}\big) - \log(S_{ot}) - (\bar{\alpha} P_{jt} + \bar{\gamma} C_j + X_{jt}^{v}\,\beta^{v} + X_j^{inv} \beta^{inv}) \tag{9}$$

In the third and final step, I apply the mean coefficient estimates ($\bar{\alpha}, \bar{\gamma}, \beta^{v}, \beta^{inv}$) to the simulated parameters (described above) to calculate consumer-product pairwise utilities, following Equation (1). Individual optimal choices and the number of purchases of each product are also directly obtained from this utility matrix, yielding product market shares. Utility and market shares of the original and treated samples are calculated based on the same group of 100,000 simulated consumers.

4.2 Results and Insights into Platform Strategies – Identification I

The results from the counterfactual experiments based on Identification I are reported in Table 13, for both product categories — i.e., men's fragrances and cell-phone wireless chargers. For each type of product, I report the number of likely authentic purchases, number of likely counterfeit purchases, number of choices of the outside option, and the average revenue per transaction.[4] The results for the Control group are treated as the

[4] A "transaction" could cover the purchase of a product or choice of outside option.

baseline and report the percentage deviations in the above variables under each of the three treatments.

Table 13. Results from the Counterfactual Experiments – Identification I

		Likely Authentic Purchases	Likely Counterfeit Purchases	Outside Option Purchases	Average Revenue per Decision
Men's Fragrances	*Control*	55,916	278,324	665,760	$10.20
	Treatment I: *Counterfeit Shock*	-8.23%	-0.78%	+1.02%	-2.04%
	Treatment II: *Counterfeit Detection*	+12.23%	+1.47%	-1.64%	+3.87%
	Treatment III: *Counterfeit Ban*	+93.56%	-100%	+33.95%	-66.22%
Cell-Phone Wireless Chargers	*Control*	507,555	756,407	436,038	$26.31
	Treatment I: *Counterfeit Shock*	-27.67%	+3.90%	+25.45%	-7.49%
	Treatment II: *Counterfeit Detection*	+25.17%	-16.29%	-1.04%	-0.27%
	Treatment III: *Counterfeit Ban*	+84.32%	-100%	+75.32%	-31.06%

Notes: This table reports the change in the percentages of likely authentic purchases, likely counterfeit purchases, number of outside option choices, and average revenue per decision, respectively, relative to the control group in three counterfactual experiments. In Treatment I (*Counterfeit Shock*) we shift the distribution of counterfeit probability in terms of stochastic dominance, so that the perceived probability of authentic product is uniformly reduced by 20%. In Treatment II (*Counterfeit Detection*) we enhance the contrast between counterfeit and authentic products by increasing counterfeit probability by 50% if it is larger than 0.5, and by reducing it by 50% if it is less than 0.5. In Treatment III (*Counterfeit Ban*) we eliminate likely counterfeit products from the market. 100,000 consumers are simulated for 10 markets (i.e., total 1,000,000 consumer decisions) for men's fragrances and 100,000 consumers are simulated for 17 markets (i.e., total 1,700,000 consumer decisions) for cell-phone wireless chargers.

Starting with the fragrance product category, under Treatment I (*Counterfeit Shock*), where there is a proliferation of counterfeit activity, we can see that product purchases (mostly likely authentic products) are lost to the outside option, and the average revenue the platform receives from one consumer transaction or decision declines by 2.04%. Under Treatment II (*Counterfeit Detection*), where counterfeiting is proactively detected by the platform and reported or signaled to consumers, we see that the share of likely authentic purchases increases by 12.23%, while non-purchases (i.e., choice of outside option) goes down, in a way that increases average revenue per transaction by 3.87%.[5] Finally, under Treatment III (*Counterfeit Ban*), purchases of likely counterfeit products are replaced by sharply increased purchases of likely authentic products — but also in an increased frequency of non-purchases (i.e., choice of outside option). The net result is that the average revenue per transaction is sharply reduced by 31.06%. The results are similar for wireless chargers, though the magnitudes of the effects under corresponding treatments are different.

We can glean a number of qualitative insights from the counterfactual experiments. First of all, if counterfeiting activity is not controlled, and deceptive knockoffs continue to proliferate in the marketplace, then this will substantially and negatively impact the welfare of authentic sellers and the platform alike. However, the extreme measure of banning all likely counterfeit listings will benefit authentic sellers for sure, but at a large cost to the platform in terms of lost revenues. The middle ground of the platform providing an enhanced counterfeit detection capability to consumers seems to align the welfare of

[5] The small increase in the proportion of likely counterfeit purchases is likely due to the noisy estimation of the random coefficient on counterfeit probability.

both authentic sellers and the platform. This compromise strategy of providing a counterfeit detection mechanism and sharing the results with consumers, rather than eliminating all suspicious listings immediately, can both reduce the damage to original sellers and maintain a satisfactory user base. Further discussion of the implications is provided in the following section.

4.3 Results and Insights into Platform Strategies – Identification II

Table 14 reports the results of counterfactual experiments based on Identification II, for men's fragrances category. To keep consistency with the experiments applied on Identification I, I categorize the listings into two groups based on the counterfeit intensity variable. Specifically, I define listings with D1 (i.e. the standardized percentage drop of the number of optional offers from Period I to Period III) no less than 50% as likely counterfeit listings, and the rest as likely authentic. Among 770 listings in the sample for Identification II, 145 of them are grouped as likely counterfeit listings. Compared to the classification in Identification I, less samples are considered as knockoffs under a less strict detection approach. This could cause the market dynamics under the three treatments to deviate to some degree from the results in Section 4.2.

The first row shows the total number of likely authentic purchases, the total number of likely counterfeit purchases, the total number of choices of the outside option across twelve markets, and the average revenue the platform receives from each consumer transaction or decision. I also report the average utility consumers gained from their purchasing decisions in the last column. Again, the results of the Control group are treated as the baseline;

following rows report the percentage deviations in the above outcomes under each of the three treatments.

Under Treatment I when there is a Counterfeit Shock, namely a proliferation of online counterfeit activities without platform regulation, the purchase of authentic listings decreases by 8.87% compared to the Control Group, and customers who leave the market without any purchase increases by 10.21%. As a result, the platform loses 8.22% of its revenue. Treatment III shows the results when the platform strictly detects and bans all suspicious listings, which could take up about 19% of all the competitors in this market. Although purchases from likely authentic sellers increases by 22.23%, with 14.85% more consumers leaving the market, the platform revenue declines by 12.33%, which is even higher than the case of Treatment I. Applying Treatment II, when the platform adopts a mild strategy to fight against knockoffs and protect its users, likely authentic sellers only faces 0.02% drop of sales and the platform gains 3.16% revenues instead. Combining sales volume with dynamic product prices, we can calculate revenues each listing gain under each market and summarize the changes in authentic seller revenues under the three cases. Results are consistent with the change of purchases. On the consumer side, consumers' mean utility decreases by 2.98% under Treatment I and 12.05% under Treatment III but increases by 4.92% under Treatment II. As consumer welfare is in align with platform revenues, an optimal strategy should be at a middle point to well balance authentic sellers' profits and the platform revenue. Although the magnitudes of experiment results diverge from the results based on Identification one, partly due to the rules and measures we follow to classify likely counterfeit and likely authentic listings, the effects are largely similar, and conclusions are consistent.

Table 14. Results from the Counterfactual Experiments – Identification II

		Likely Authentic Purchases	Likely Counterfeit Purchases	Outside Option Purchases	Average Revenue per Decision	Average Consumer Utility
Men's Fragrances	*Control*	516,754	188,349	494,897	$18.79	1.55
	Treatment I: *Counterfeit Shock*	-8.87%	-2.50%	+10.21%	-8.22%	-2.98%
	Treatment II: *Counterfeit Detection*	-0.02%	+13.57%	-5.14%	+3.16%	+4.92%
	Treatment III: *Counterfeit Ban*	+22.23%	-100%	+14.85%	-12.33%	-12.05%

Notes: This table reports the change in the percentages of likely authentic purchases, likely counterfeit purchases, number of outside option choices, and average revenue per decision, respectively, relative to the control group in three counterfactual experiments. In Treatment I (*Counterfeit Shock*) we shift the distribution of counterfeit intensity in terms of stochastic dominance, so that the perceived probability of authentic product is uniformly reduced by 20%. In Treatment II (*Counterfeit Detection*) we enhance the contrast between counterfeit and authentic products by increasing counterfeit intensity by 50% if it is larger than 0.5, and by reducing it by 50% if it is less than 0.5. In Treatment III (*Counterfeit Ban*) we eliminate likely counterfeit products from the market. 100,000 consumers are simulated for 10 markets (i.e., total 1,000,000 consumer decisions) for men's fragrances.

CHAPTER 5

Summary and Conclusions

5.1 Conclusions and Managerial Implications

Online retail markets are increasingly plagued by the presence of deceptive counterfeit products, a problem that is exacerbated on Amazon by the proliferation of third-party sellers on the platform. Indeed, many brands have left Amazon, such as Nike and PopSockets (Barkho 2020). Despite the growing recognition of the online counterfeiting problem, there is very little research that has addressed it, largely due to the challenge of identifying counterfeit products. Applying text mining techniques to review data, leveraging an exogenous event with original manufacturers' knowledge and the platform's technology combined, in conjunction with a BLP-style random choice model, I have developed a comprehensive empirical framework to *identify* deceptive online counterfeiting, modeled it into consumers' decision-making process, and assess its economic *impact* on authentic sellers and the platform alike. I leverage the structural parameter estimates to conduct a number of counterfactual experiments to generate further insights into the impacts of counterfeiting, and how it could be countered and mitigated.

To summarize the findings, this work shows that it is indeed feasible to leverage product reviews to identify counterfeiting activity in Amazon listings, through suitable machine learning models — the best model has achieved 83% accuracy, a level that can be further enhanced by enlarging the training data set. The detection results are also confirmed by platform artificial intelligence algorithms automatically scanning and protecting the marketplace. Clear evidence is found that counterfeiting negatively impacts user utility, especially for high-end products, best sellers, and utilitarian or functional products. Products that are likely counterfeits are net substitutes for products that are likely authentic, and counterfeit sellers appear to take away significant market share from authentic sellers — in the aggregate — especially for luxury and expensive products. The counterfactual experiments shed further light on the pernicious effect of counterfeiting on market shares of authentic sellers and platform revenues. Yet, simply eliminating listings with counterfeiting activity is like "throwing the baby with the bath water" in that a strict ban of listings with counterfeiting activity would significantly hamper platform revenues. The answer lies in deploying counterfeit detection algorithms and sharing the results with consumers. My research finds that this approach aligns the interests of authentic sellers, consumers, and that of the platform.

This work provides implications for consumers, sellers and the platform. For consumers, this study confirms the potential of user-generated reviews in identifying deceptive counterfeits, paying most attention to reviews with helpful votes and ones that explicitly mention counterfeiting — even if they are buried deep in the list of reviews. This extra effort is particularly rewarding for buyers of luxury and high-end products, where we find evidence of disproportionate counterfeiting activity. The percentage of four-star ratings are often ignored, in lieu of the attention focused on one-star and five-star reviews, but we found them to be predictive of product authenticity.

For sellers of authentic products, the results suggest two key implications. The first has to do with the negative spillover from counterfeit sellers that join the same ASIN listing. A strategy to deal with this follows from theories of quality signaling (e.g., Milgrom and Roberts 1986), wherein high-quality sellers should provide costly signals of quality that low-quality sellers cannot afford, such as product return warranty or genuineness certification. This type of costly signaling could be enabled by a unique code attached to each authentic product by which consumers can validate the authenticity of the product or initiate a product return. Second, authentic sellers should monitor their product review pages to track signs markers of counterfeiting activity — especially disclosing reviews — and they should be prepared to directly respond to messages in the review system to allay concerns about their product quality. They may also consider leaving the listing or the platform altogether if the negative impacts from counterfeiting are significant enough.

From the perspective of platform operators, the counterfeiting problem is a vexing one to deal with. On the one hand, the entry of counterfeiters objectively increases market size on the seller side and subsequently the consumer side, through supply-side economies of scale. At the same time counterfeit sellers steal market shares from authentic sellers, there is a danger that counterfeiting could drive authentic sellers away from the market, with dire consequences for the traded quality of products in the market. Yet, as the counterfactual experiments indicate that banishing listings with credible counterfeit activity, while enhancing the market shares of authentic sellers, would come at a prohibitive reduction in revenues of the platform. A more productive direction is for the platform to deploy counterfeit detection algorithms — perhaps like the ones that we develop, based on review texts — and report the results to the users through explicit counterfeit markers added to each ASIN listing. The counterfactual experiments indicate that such an approach would benefit authentic sellers and the platform alike. The platform needs to strike the right balance between dissuading counterfeiting activity through information disclosure and supporting sales overall sales volume on the platform, wherein the presence of comingled counterfeit sellers provides a sort of liquidity effect.

5.2 Future Work

This work is not without limitations. First, both of the identification approaches are at the listing rather than seller level, due to the fact that it is common for multiple sellers to offer a product under the same ASIN listing, many of whom are third-party sellers. This is particularly true for horizontal taste products, rather than vertical functional or utilitarian products. Relatedly, user reviews do not indicate which seller the purchase came from. Second, in Identification I there is no objective ground truth for counterfeiting activity. Rather, our training data set is constructed by human coders that label listings based on their subjective assessments of a likely counterfeit encounter when purchasing from an ASIN listing. I take measures to try to counteract the endogeneity of counterfeit probability, but there is no doubt some noise in the estimation of counterfeit probability. Having said that, the platform is probably in the best position to monitor and characterize counterfeiting activity by leveraging transaction data that only they are privy to. Therefore, I leverage Project Zero to develop Identification II to directly measure the intensity of existing counterfeit activities before this service was largely provided to original brands.

A third limitation lies in the potential frauds and lack of timeliness in online reviews. Disclosing reviews can be diluted by fake positive reviews, which may bias the measure of counterfeit intensity. It is also not clear if related reviews from previous purchases will be removed when Project Zero detects and removes suspicious offers. These questions could both limit the efficiency of the review-related identification. Another limitation is that I only study two product categories because of the painstaking efforts required to compile training data sets, and to specify detailed product characteristics. I have selected representative horizontal and vertical taste products, but there is clearly room for future studies to examine other types of products or focus on a platform-wide characterization of where counterfeiting is most likely to be detected.

In addition, I design three counterfactual experiments to study platform strategies in reaction to online counterfeit trades and compare their pros and cons. However, authentic sellers themselves can also respond to this issue and change the dynamic of the online market. For example, they can slightly adjust their prices by distributing some coupons, to gain more purchases in the competition. They can also register multiple "distributor" stores to lower the likelihood of consumers unconsciously selecting a counterfeit third-party seller. Adopting the serialization service provided by Amazon is another approach. By attaching a unique code, they can signal their authenticity to the platform and consumers. Future research can design corresponding counterfactual experiments by modifying authentic product prices, unobserved product characteristics, and the number of authentic offers to observe the effect on authentic sellers themselves, consumer utility, and the platform profits.

Lastly, the main effect of counterfeiting on consumer utility is captured, but the choice model doesn't account for spillover effects at a platform level such as customer churn caused by the poor efficiency of matching in a thick market or the loss of consumer trust. Again, these are fruitful directions for future research. Despite these limitations, this is the first comprehensive empirical study of identification and impact of online deceptive counterfeit products.

REFERENCES

Archak N, Ghose A, Ipeirotis PG (2011) Deriving the pricing power of product features by mining consumer reviews. *Mgmt. Sci*. *57*(8): 1485-1509.

Amazon (2023), Amazon Project Zero – Helping drive counterfeits to zero, *Amazon*, https://brandservices.amazon.com/projectzero.

Bai X, Marsden JR, Ross WT Jr, Wang G (2020) A note on the impact of daily deals on local retailers' online reputation: mediation effects of the consumer experience. *Info. Sys. Res.* 31(4): 1132-1143.

Barkho G (2020) Why some DTC brands are pulling out of Amazon. *Modern Retail* (February 5), https://www.modernretail.co/startups/dtc-brands-leaving-amazon/.

Baum CF, Schaffer ME, Stillman S (2003) Instrumental variables and GMM: Estimation and testing. *The Stata Journal.* 3(1): 1-31.

Berry S, Levinsohn J, Pakes A (1995) Automobile prices in market equilibrium. *Econometrica: Journal of the Econometric Society*. 841-890.

Berry S, Pakes A (2007) The pure characteristics demand model. *International Economic Review*. 48(4): 1193-1225.

Berzon A, Shifflett S, Scheck J (2019) Amazon Has Ceded Control of Its Site. The Result: Thousands of Banned Unsafe or Mislabeled Products, *The Wall Street Journal* (August 23), https://www.wsj.com/articles/amazon-has-ceded-control-of-its-site-the-result-thousands-of-banned-unsafe-or-mislabeled-products-11566564990.

Bressler MS, Bressler LA (2018) Makers vs. Fakers: how counterfeit goods hurt competition, harm the economy, and kills consumers. *available from Researchgate.net*.

Chaloux JM, Hayes JP, Aggarwal M, Tsatsoulis PD, Accenture Global Solutions Ltd (2020) Detection of counterfeit items based on machine learning and analysis of visual and textual data. *U.S. Patent* 10,691,922.

Chevalier J, Goolsbee A (2003) Measuring prices and price competition online: Amazon. com and BarnesandNoble. com. *Quantitative marketing and Economics*. 1(2): 203-222.

Cheung M, She J, Liu L (2018). Deep learning-based online counterfeit-seller detection. *IEEE INFOCOM 2018-IEEE Conference on Computer Communications Workshops (INFOCOM WKSHPS).* (April): 51-56.

Chevalier JA, Mayzlin D (2006) The effect of word of mouth on sales: Online book reviews. *Journal of marketing research*. 43(3): 345-354.

Cho WY, Ahn BH (2010) Versioning of information goods under the threat of piracy. *Information Economics and Policy*. 22(4): 332-340.

Conlon C, Gortmaker J (2020) Best practices for differentiated products demand estimation with pyblp. *The RAND Journal of Economics*. 51(4): 1108-1161.

Coppola D (2022) Quarterly share of e-commerce sales of total U.S. retail sales from 1st quarter 2010 to 1st quarter 2022. *Statista* (August 7), https://www.statista.com/statistics/187439/share-of-e-commerce-sales-in-total-us-retail-sales-in-2010/.

Decker R, Trusov M (2010) Estimating aggregate consumer preferences from online product reviews. *International Journal of Research in Marketing*: 27(4): 293-307.

Dewan S, Hsu V (2004) Adverse selection in electronic markets: evidence from online stamp auctions. *Journal of Industrial Economics. LII (4)*: 497-516.

Dimoka A, Hong Y, Pavlou PA (2012) On product uncertainty in online markets: Theory and evidence. *MIS quarterly*. 395-426.

Droesch B (2021) Amazon dominate US ecommerce, though its market share varies by category, *eMarketer* (April 27) https://www.emarketer.com/content/amazon-dominates-us-ecommerce-though-its-market-share-varies-by-category.

Ebay (2023), Counterfeit Item Policy, *Ebay*, https://www.ebay.com/help/policies/prohibited-restricted-items/counterfeit-item-policy?id=4276.

Gallagher RJ, Reing K, Kale D, Ver Steeg G (2017) Anchored correlation explanation: Topic modeling with minimal domain knowledge. *Transactions of the Association for Computational Linguistics*. 5: 529-542.

Ghose A, Han SP (2014) Estimating demand for mobile applications in the new economy. *Management Science*. 60(6): 1470-1488.

Ghose A, Ipeirotis PG, Li B (2012) Designing ranking systems for hotels on travel search engines by mining user-generated and crowdsourced content. *Marketing Science*. 31(3): 493-520.

Ghose A, Sundararajan A (2006) Evaluating pricing strategy using e-commerce data: Evidence and estimation challenges. *Statistical Science*. 21(2): 131-142.

Givon M, Mahajan V, Muller E (1995) Software piracy: Estimation of lost sales and the impact on software diffusion. *Journal of Marketing*. 59(1): 29-37.

Hernandez B (2021) What Is Amazon Project Zero and How It Helps Fight Counterfeits, *Perpetua*, https://perpetua.io/blog-amazon-project-zero/.

Hui KL, Png I (2003) Piracy and the legitimate demand for recorded music. *Contributions to Economic Analysis and Policy*. 2(1): 1-22.

Jain S (2008) Digital piracy: A competitive analysis. *Marketing science*. 27(4): 610-626.

Kennedy JP (2020) Counterfeit Products Online. Holt TJ, Bossler AM, eds. *The Palgrave Handbook of Cybercrime and Cyberdeviance*, 1001-1024.

Lahiri A, Dey D (2013) Effects of piracy on quality of information goods. *Management Science*. 59(1): 245-264.

Lewis G, Zervas G (2016) The welfare impact of consumer reviews: A case study of the hotel industry. Working paper, UPenn, Philadelphia.

LocalCircles (2018), Counterfeit or fake products on eCommerce sites is a much bigger problem than we thought, *LocalCircles*, https://www.localcircles.com/a/press/page/fake-products-on-ecommerce-sites.

Lu S, Wang X, Bendle N (2020) Does Piracy Create Online Word of Mouth? An Empirical Analysis in the Movie Industry. *Management Science*. 66(5): 2140-2162.

McFadden D (1973) *Conditional logit analysis of qualitative choice behavior.*

Mehta D (2020), Amazon Project Zero launches in 7 new countries, *Amazon*, https://www.aboutamazon.com/news/company-news/amazon-project-zero-launches-in-7-new-countries.

Milgrom P, Roberts J (1986). Price and advertising signals of product quality. Journal of Political Economy. 94(4): 796-821.

Pieterse J (2021) Are Fragrances On Amazon Real?, *Fragrancetoday*, https://fragrancetoday.com/are-fragrances-on-amazon-real/.

Qian Y (2014) Counterfeiters: Foes or friends? How counterfeits affect sales by product quality tier. *Management Science*. 60(10): 2381-2400.

Qian Y, Gong Q, Chen Y (2015) Untangling searchable and experiential quality responses to counterfeits. *Marketing Science*. 34(4): 522-538.

Quora (2019) How much perfume sold on Amazon is counterfeit?. Retrieved September 7, 2022, https://www.quora.com/How-much-perfume-sold-on-Amazon-is-counterfeit.

Sharma A, Srinivasan V, Kanchan V, Subramanian L (2017) August. The fake vs real goods problem: microscopy and machine learning to the rescue. *The 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2011-2019.

Silcox K (2021) The Hidden Costs of Buying Fake Cosmetics on Sites Like Amazon, *Money* (January 19), https://money.com/fake-amazon-makeup-skincare/.

Smith MD, Telang R (2009) Competing with free: The impact of movie broadcasts on DVD sales and Internet piracy. *MIS Quarterly*. 321-338.

Song M (2015) A hybrid discrete choice model of differentiated product demand with an application to personal computers. *International Economic Review*. 56(1): 265-301.

Steele C (2019) The Ugly Truth About Amazon Beauty, *PCMAG* (September 17), https://www.pcmag.com/opinions/the-ugly-truth-about-amazon-beauty.

Sundararajan, A., 2004. Managing digital piracy: Pricing and protection. *Information Systems Research*. 15(3): 287-308.

Vincent DW (2015) The Berry–Levinsohn–Pakes estimator of the random-coefficients logit demand model. *The Stata Journal*. 15(3): 854-880.

Wang Q, Li B, Singh PV (2018) Copycats vs. original mobile apps: A machine learning copycat-detection method and empirical analysis. *Information Systems Research*. 29(2): 273-291.

Zeibak (2020) How to Win the Amazon Buy Box in 2021, *Amazon & Marketplaces* (March 25), https://tinuiti.com/blog/amazon/win-amazon-buy-box/.