

ABSTRACT

Title of Dissertation: EMPIRICAL INVESTIGATION OF USERS’
SUCCESSFUL STRATEGIES IN ONLINE
PLATFORMS - EVIDENCE FROM CROWD-
SOURCING AND SOCIAL MEDIA
PLATFORMS

Mikhail Viktorovich Lysyakov,
Doctor of Philosophy, 2021

Dissertation directed by: Professor Siva Viswanathan,
Department of Decision, Operations and
Information Technologies

With the proliferation and constant growth of online platforms, there has been an increasing interest among academicians and practitioners to understand various aspects of these platforms, including the effective design of platforms, their governance and user engagement. This dissertation seeks to add to this stream of research by leveraging large-scale unstructured data and corresponding data analytics and econometric techniques to examine users’ strategies in online social media and crowdsourcing platforms and gain insights into factors that lead to successful outcomes.

The first essay examines the content strategies of closely competing firms on Twitter with a focus on how the similarity/dissimilarity of their content strategies impacts their online outcomes. I find that firms that are more adept at leveraging higher-level social media affordances, such as interactivity, collaboration, and online contests

to differentiate their content strategies experience better outcomes as compared to their closest rivals that only leverage the basic technological affordances of social media.

The second essay examines successful strategies of users (designers) in a crowdsourcing platform wherein clients post contests to solicit design solutions for a monetary reward. This study uses state-of-the-art deep learning and image analysis techniques to examine the strategies of experienced and less-experienced designers in open contests where later-entrants can potentially leverage information spillovers from earlier design submissions within a contest. I find that while later-entrants typically leverage information spillovers from earlier submissions in a contest, only experienced designers who are able to integrate information from multiple highly-rated early submissions are more likely to be successful.

The third essay examines users' strategies in response to the introduction of an Artificial Intelligence system for logo design in an online crowdsourcing design platform. In analyzing what differentiates successful contestants from the others, I find that the successful contestants significantly increase focus (i.e., the number of re-submissions per contest) and increase the emotional content as well as the complexity of their designs, in response to the introduction of the AI system.

Collectively, the findings from these studies add to our understanding of successful strategies in online platforms and provide valuable insights to theory and practice.

EMPIRICAL INVESTIGATION OF USERS' SUCCESSFUL STRATEGIES IN
ONLINE PLATFORMS - EVIDENCE FROM CROWD-SOURCING AND
SOCIAL MEDIA PLATFORMS

by

Mikhail Viktorovich Lysyakov

Dissertation submitted to the Faculty of the Graduate School of the
University of Maryland, College Park, in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
2021

Advisory Committee:

Professor Siva Viswanathan, Chair

Professor Guodong "Gordon" Gao

Professor Il-Horn Hann

Professor Kunpeng Zhang

Professor Anna Alberini (Dean's Representative)

© Copyright by
Mikhail Viktorovich Lysyakov
2021

Acknowledgements

First of all, I would like to acknowledge the tremendous support of my advisor Professor Siva Viswanathan, the Chair of my dissertation committee. Professor Siva has constantly challenged me to improve every aspect of my research program, and directed me to become a fully independent researcher, which I greatly appreciate. He guided me through the whole PhD program up to a very successful outcome – a Tenure Track position in academia. I also want to thank Professors in my dissertation committee for their very helpful guidance and feedback – Professor Il-Horn Hann, Professor Guodong “Gordon” Gao, Professor Kunpeng Zhang and Professor Anna Alberini. I also want to thank Professor Anandasivam Gopal for his guidance during my initial years in the PhD program.

My PhD experience has been greatly enhanced by frequent interactions with my top-performing peers, especially Raveesh Mayya, Weiguang Wang, Sabari Karmegam, Che-Wei Liu, Aishwarya Shukla, Lanfei Shi, Salman Aljazzaf, Jaehoon Jeong, Abdullah Alhauili, Lauren Dahlin and Vidit Mehta. I would also like to acknowledge the great support from our Director of Doctoral Programs Justina Blanco.

Finally, I would like to thank my wife, Yuliya, an amazing partner, for her very important role and support during the whole PhD program experience.

Table of Contents

Acknowledgements.....	ii
Table of Contents.....	iii
List of Tables	v
List of Figures	vi
Chapter 1: Introduction and Overview	1
Chapter 2: Retailers’ Content Strategies on Social Media: Insights from Analysis of Large-scale Twitter Data.....	7
2.1 Introduction.....	7
2.2 Related Research and Theoretical Background	12
2.2.1 Social Media Content Categories.....	14
2.2.2 Social Media Affordances.....	17
2.3 Research Context and Data.....	20
2.4 Methodology	22
2.4.1 Similarity/Dissimilarity.....	22
2.4.2 Followers.....	25
2.5 Hypotheses and Econometric Specification.....	26
2.5.1 Hypotheses	26
2.5.2 Econometric specification.....	28
2.6 Results.....	32
2.7 Why Do Dissimilar Firms Perform Better?	36
2.8 Implications and Conclusion.....	45
Chapter 3: Synthesizing Winning Strategies: What Differentiates Experienced Designers in Crowdsourcing Markets?.....	51
3.1 Introduction.....	51
3.2 Related Research and Theoretical Underpinnings.....	57
3.2.1 Design of Contests	57
3.2.2 Designers’ Strategies	58
3.2.3 Information Spillovers	59
3.2.4 Experienced Designers and Recombinant Innovations.....	62
3.2.5 Hypotheses.....	64
3.3 Research Context	67
3.3.1 Contest Description.....	67
3.4 Data and Methodology.....	68
3.4.1 Structural Similarity.....	68
3.4.2 Empirical Models.....	69
3.5 Results.....	71
3.5.1 Effect of Information Provided by the Contest Holder.....	71
3.5.2 Effect of Information Spillovers from Prior Submissions	72
3.5.3 Effects of Timing of Entry and Designer Experience.....	73
3.5.4 How do Experienced Designers Differ from Other Designers?	74
3.5.5 Comparison of Open and “Blind” Contests	79
3.5.6 Robustness Check for a Potential Confounding Effect of Task Textual Descriptions	80

3.6 Implications and Conclusion.....	81
Chapter 4: Threatened by AI: Analyzing Users' Responses to the Introduction of AI in a Crowd-sourcing Platform.....	85
4.1 Introduction.....	85
4.2 Related Research.....	92
4.2.1 Impacts of AI on Organizations.....	92
4.2.2 Contestants' Behaviors and Strategies on Crowdsourcing Platforms.....	94
4.2.3 AI Systems and Prior Technologies.....	95
4.2.4 Emotions and Complexity.....	98
4.3 Research Context and Data.....	99
4.4 Methodology.....	101
4.4.1 Task Descriptions' Requirements Classification.....	101
4.4.2 Complexity of Design Images.....	103
4.4.3 Emotions in Design Images.....	104
4.4.4 Relationship between Emotions and Complexity.....	106
4.5 Results.....	106
4.5.1 Effects of AI on the Number and Composition of Contests.....	106
4.5.2 Designers' Responses to the Launch of the AI system for Logo Design.....	108
4.5.3 Successful Designers' Responses to the Introduction of the AI System for Logo Design.....	110
4.5.4 Additional Controls.....	112
4.5.5 Difference-in-differences Analysis with Propensity Score Matching.....	114
4.5.6 Robustness Checks.....	117
4.6 Implications and Conclusion.....	118
Chapter 5: Conclusion.....	123
Tables.....	129
Figures.....	143
Appendices.....	148
Bibliography.....	171

List of Tables

Table 2.1. Descriptive statistics of the dataset	129
Table 2.2. Descriptive statistics for the main variables	129
Table 2.3. Fixed effects, GMM and 2SLS similarity model results for the period Jan. 2012-Aug. 2017	130
Table 2.4. Beta regression model with a logit link results for each higher-level category	130
Table 2.5. Normalized engagement (number of (favorites/retweets)/tweet/100,000 followers) by category	131
Table 2.6. Two-stages least squares model results for the effect of usage of higher-level categories on online engagement and new followers' acquisition rate	131
Table 3.1. Descriptive statistics for the main variables	132
Table 3.2. Effect of structural similarities on the probability of winning in models with contest fixed effects	134
Table 3.3. Timing of entry by experience in models with contest fixed effects	135
Table 3.4. Effects of experience in open and "blind" contests in models with contest fixed effects	135
Table 4.1. Descriptive statistics for the main variables in the dataset	136
Table 4.2. Confusion matrix for 5 emotions	136
Table 4.3. Descriptive statistics for the variables related to emotions and complexity	137
Table 4.4. Comparison of lower-tier and higher-tier logo contest task descriptions before and after the launch of AI	138
Table 4.5. Propensity score matching model for designers' switching behaviors after the AI launch	138
Table 4.6. Comparison of the number of contests that each designer participates in per day before and after AI for 3 groups of successful designers as compared to the unsuccessful designers	139
Table 4.7. Comparison of the number of re-submissions per user per contest before and after AI for 3 groups of successful designers as compared to the unsuccessful designers	139
Table 4.8. Comparison of presence of emotional content before and after the AI launch for 3 groups of successful designers	140
Table 4.9. Comparison of complexity before and after the AI launch for 3 groups of successful designers	140
Table 4.10. Comparison of five emotions before and after the AI launch for 3 groups of successful designers	141
Table 4.11. Comparison of presence of emotional content before and after the AI launch for 3 groups of unsuccessful designers	141
Table 4.12. Comparison of complexity before and after the AI launch for 3 groups of unsuccessful designers	142
Table 4.13. Comparison of five emotions before and after the AI launch for 3 groups of unsuccessful designers	142

List of Figures

Figure 2.1. Content similarity with top competitors	143
Figure 2.2. Examples of top competitors with one outlier (Panel A) and two outliers (Panel B)	143
Figure 2.3. Hierarchy of firm-tweet categories.....	144
Figure 3.1. Structural similarity index example of similar images (similarity is 0.92)	145
Figure 3.2. Structural similarity index example of dissimilar images (similarity is 0.14)	145
Figure 3.3. Two original highly-rated images (A and B) and their synthesized images (C and D).....	146
Figure 3.4. Source images A and B, synthesized images and corresponding winning images	146
Figure 4.1. Examples of a simple logo (Image A, complexity score: 0.235) and a complex logo (Image B, complexity score: 2.05).....	147
Figure 4.2. Examples of logos with emotional content (image C) and without emotional content (image D)	147
Figure 4.3. The number of lower-tier and higher-tier logo contests before and after AI	147

Chapter 1: Introduction and Overview

There are a plethora of online platforms and they are constantly growing. Some examples of these platforms are social media platforms that include platforms for forums and message boards (Reddit), social networks (Facebook), review and opinion sites (Yelp), blogging and microblogging (Twitter), media sharing (YouTube). Other examples are crowdsourcing platforms such as platforms for open innovation (InnoCentive), crowdfunding (Kickstarter), crowdttesting (test IO), collaborative knowledge (openIDEO), and microtasking (Amazon Mechanical Turk).

With the proliferation and rapid growth of online platforms (de Reuver et al. 2018), there has been an increasing interest among academicians and practitioners to understand various aspects of these platforms, such as an effective design of platforms, their governance and user engagement (Constantinides et al. 2018). Prior research has studied different online platforms, such as social media and crowd-based platforms, but there is still a lack of understanding of the strategies that lead to successful outcomes for users on these platforms.

For instance, while social media research has studied extensively interactions among users (Kietzmann et al. 2011) as well as how users produce and exchange user-generated content (Susarla et al. 2011, Smith et al. 2012) and react to content and word-of-mouth produced by other users (Adamopoulos et al. 2015), less attention has been devoted to firms' social media content strategies and to how those content strategies may help achieve higher engagement and higher growth in userbase in social media brand communities (Bapna et al. 2019). Similarly, research on crowdsourcing platforms has examined the optimal design of contests on such platforms and user

behaviors (Segev 2020), including the optimal prize structure (Archak and Sundararajan 2009), the impact of “open” and “blind” contests (Jian et al. 2017), as well as reaction of users to competition (Gross 2016) and to the presence of “superstar” users (Zhang et al. 2019). However, there is hardly any research that has examined successful strategies employed by users in online crowdsourcing contests. Access to large-scale data and techniques for analyzing unstructured data enables us to gain insights into users’ strategies in such platforms. This dissertation leverages large-scale data and corresponding data analytics and econometric methods to gain insights into factors that drive successful strategies for users in social media and crowdsourcing platforms.

The first essay of the dissertation seeks to understand firms’ successful content strategies on social media and to gain insights into the role of those strategies in helping firms achieve higher online engagement and grow their online brand communities of followers. While recent research has started examining firms’ content strategies on social media (Lee et al. 2018, Bapna et al. 2019), there is still a gap in this research stream relating to competitors’ content strategies on social media and the impact of those content strategies on related outcomes. This study seeks to fill the gap by identifying similar and dissimilar content strategies of traditional close competitors on Twitter. While close competitors have a competitive pressure to be isomorphic, i.e., more similar, in their Web footprints (Pant and Sheng 2015), I find that some firms are dissimilar from their close rivals in their Twitter content strategies, while other firms are similar to their close competitors in their content strategies. I further explore whether similarity or dissimilarity in content strategies leads to better online outcomes

and find that dissimilar firms have better social media outcomes compared to their more similar rivals. To understand the underlying mechanism, I use deep learning models to classify all firm-tweets into 10 categories that correspond to social media affordances identified in prior research (Karahanna et al. 2018). Leveraging this classification, I propose a three-tier framework of content strategies and corresponding affordances wherein the lowest “content” tier includes only basic affordances of communication and self-presentation, while the middle “community” tier includes affordances of meta-voicing and relationship formation, and, finally, the top “co-creation” tier corresponds to affordances of interactivity, collaboration and competition. I find that dissimilar firms are more likely to use content strategies in the higher “community” and “co-creation” tiers as compared to their more similar rivals. I further find that leveraging content strategies in higher tiers leads to higher online engagement on social media and attracts more new followers to the focal brand community. These findings contribute to research on competitors’ successful content strategies and social media affordances, and also have interesting practical implications for firms’ content strategies on social media platforms.

The second essay seeks to understand users’ successful strategies on a crowdsourcing platform for design tasks wherein clients (contest holders) initiate contests to solicit solutions from participants (designers) for a monetary reward. Prior research in crowdsourcing has found that experienced designers are more likely to win in contests compared to less-experienced designers (Khasraghi and Aghaie 2014). However, what is it that experienced designers do that gives them an edge in such contests has not been examined. I specifically focus on open contests for logo design

where information and feedback (star rating) from prior logo image submissions is available to other designers who enter contests later. Thus, in these open contests designers who enter a contest later can potentially leverage information spillovers from earlier designers' submissions. On the one hand, those spillovers might help later-entering designers to learn about a contest holder's preferences and increase their probability of success. On the other hand, there is a higher likelihood of direct imitation of prior submissions and contest holders might react negatively to such direct imitation. I leverage this setting and employ state-of-the-art image analysis techniques to directly measure information spillovers from prior highly-rated submissions and to understand how experienced designers are different from less-experienced designers in leveraging such spillovers. I find that experienced designers are less likely to excessively copy each individual prior highly-rated submission compared to less-experienced designers. Interestingly, experienced designers are more likely to synthesize information from several prior highly-rated submissions as compared to less-experienced designers, and it is the synthesis of such information that positively affects their probability of winning. This information synthesis is consistent with the emerging perspective on recombinant innovations, which states that most innovations happen by recombining prior innovations rather than by developing novel solutions from scratch. Thus, I find that experienced designers are more adept at recombining information spillovers from prior highly-rated designs, and that the recombination strategy gives experienced designers an edge in these contests. These findings have significant implications for both the participants in online crowdsourcing markets, as well as the designers of such marketplaces and platforms.

The third essay examines users' successful strategies in response to the introduction of an Artificial Intelligence system for logo design that can compete with human designers in a crowdsourcing design platform. This is a nascent research stream where prior research has begun to examine the impact of introduction of robots/AI systems on labor and skill composition (Frey and Osborne 2017, Dixon et al. 2019). While prior studies of the advent of AI in businesses have largely focused on the impact of AI on firms, this is one of the first studies to examine the impact of AI in a decentralized marketplace, and the heterogeneity in user responses to the threat from AI. This study seeks to understand the strategies that designers use to respond to the introduction of the AI system for logo design considering the introduction of the AI system as an external shock for the users. Majority of the contests on the platform (about 90%) are related to logo design, while the remaining contests (about 10%) are related to other categories of design such as T-shirt design. I find that after the introduction of the AI system for logo design the number of lower-tier contests for logo design, i.e., less-complex contests with a low reward amount, decreases by about 25%, while the number of higher-tier contests for logo design with a higher reward amount decreases by about 5%. In contrast, the number of contests for non-logo categories increases by 10%. Exploring designers' choice of contests after the introduction of the AI system, I find that some designers increase the proportion of higher-tier contests for logo design in which they participate while other designers increase their participation in other categories not related to logo design. Exploring heterogeneity in the contests' choice, my findings indicate that after the introduction of the AI logo system as compared to the period before the introduction, designers who have primarily participated in

contests for lower-tier, less-complex, logo designs continue to participate in these contests, while designers who have had prior exposure to more-complex logo-design contests, switch to higher-tier more-complex logo-design contests. Designers who have had prior exposure to other categories of design tasks switch to participating in contests in other categories. Interestingly, the successful strategy to respond to the AI system is to increase emotional expression and complexity of design submissions, which is the strategy adopted by the winning designers. This strategy is especially meaningful given that AI systems are limited in those dimensions of emotions and complexity. Additionally, I find that the successful designers substantially increase focus (i.e., the number of re-submissions) within each contest without changing the number of contests that they participate in, while the unsuccessful designers either decrease or do not change the number of re-submissions within each contest and participate in more contests after the AI launch. These findings contribute to the nascent research on the effects of AI on users' behaviors and strategies that help them to successfully compete with an AI system in a crowd-based platform by leveraging such human traits as emotional expression and creativity.

Collectively, the findings of the studies in this dissertation help us understand users' successful strategies and behaviors on social media and crowdsourcing platforms and also provide valuable insights for theory and practice.

Chapter 2: Retailers' Content Strategies on Social Media: Insights from Analysis of Large-scale Twitter Data

2.1 Introduction

Social media platforms (SMPs) such as Facebook, Twitter, LinkedIn, and Instagram, that primarily focus on connecting users and facilitating sharing of content among users have increasingly become important channels for firms and organizations as well. While these social media platforms have been largely used as information dissemination channels by firms, over the last few years firms have found innovative and varied uses of these platforms. For example, firms use SMPs for customer service or use SMPs to target influential social media users and bloggers, who have thousands of followers, and leverage these users as brand ambassadors to promote their products on social media platforms. Firms also use SMPs to build online communities around brands to create highly engaged audiences to promote their products and services through offline and online word of mouth. More importantly, these social media platforms have evolved into new channels for firms to compete with their rivals. Commensurate with the growth in firms' use of SMPs, there has been a growing interest in understanding how firms use SMPs. Prior research on firms use of SMPs has focused on specific content shared by firms (Swani et al. 2013, Lee et al. 2018, Bapna et al. 2019), on the community building strategies by large brands (Culnan et al. 2010), customers' complaints handling by firms (Einwiller and Steilen 2015), and offensive and defensive social media marketing strategies of firms after product-harm crises (He et al. 2017). However, there is very little research examining *how* firms' content strategies on social media platforms compare with those of their traditional rivals and what their effects are on related outcomes.

This paper seeks to fill this gap and focuses on studying how firms' content strategies on Twitter, one of the more widely and more frequently used online social media platforms, compare with those of their traditional rivals. I focus on firms that are identified as close competitors in the traditional context and examine whether these traditional competitors adopt *similar* content strategies on Twitter, and the impacts of such *similarity* or *dissimilarity* in competition on related outcomes. I also restrict my focus to B2C (business-to-consumer) firms, which operate in the retail sector with NAICS codes 44-45 since retail firms have one of the highest levels of consumer facing social media activities (adweek.com 2015). Second, I am interested in exploring traditional competitors' strategies on social media platforms, and competition in the retail sector has one of the highest growth levels among all sectors, as measured by the Herfindahl-Hirschman Index (Autor et al. 2017). The most common measures of competition are industry codes (Hoberg and Phillips 2016), and Hoover's database (Pant and Sheng 2015). These measures are extensively used by researchers in finance, strategy (Bergen and Peteraf 2002), marketing, management (Gur and Greckhamer 2018), and Information Systems (Pant and Sheng 2015). In keeping with prior research, I use the Hoover's database to identify closest competitors in the retail sector and examine the online content strategies of these competing firms. With the growing competition in the retail sector, and with retail firms' active presence on social media platforms, it is expected that offline competition will also be reflected in competition on SMPs.

While SMPs have rapidly emerged as the new frontiers of competition for traditional firms, SMPs also differ from traditional channels in important ways as they

offer a number of capabilities that provide firms the ability to adopt strategies different from offline settings. Compared to traditional offline channels, firms have an opportunity to communicate with a large number of followers simultaneously and in a personalized manner. SMPs also enable firms to share content in real time and directly monitor users' reactions to their content. SMPs also enable firms to create multiple communities of interest around specific products and services. Additionally, social media platforms allow employing a strategy of co-creating and innovating by sourcing ideas from the online "crowd", evaluating those ideas and rewarding top achievers (Mandviwalla and Watson 2014). While SMPs afford firms the ability to leverage a plethora of unique features, not all firms might be adept at leveraging the capabilities offered by SMPs and different firms might use SMPs in very different ways. Given the unique capabilities afforded by SMPs, this study seeks to examine if firms that compete closely in traditional channels also adopt similar content strategies on social media platforms. Specifically, I focus on firms' content strategies on Twitter and examine how similar or dissimilar are firms with respect to their closest traditional competitors.

Prior work has used content analysis to understand user behaviors as well as firms' behaviors in online and offline settings. A few studies have used textual analysis of online content to understand the relationships between user sentiment and emotions and brand perceptions (Ibrahim et al. 2017), as well as firm performance (Sul et al. 2014). Pant and Sheng (2015) have used textual analysis of Web footprints to show that competitors have more similar Web footprints compared to non-competing firms. In keeping with these studies, I analyze the textual content of firms' tweets to

understand how they compare with their closest traditional rivals on Twitter and the implications of such competition.

In previewing the findings, I find that close competitors that have a high degree of similarity offline show greater divergence in their content strategies online. I find that the more dissimilar a firm's content strategy is to its closest rivals, the higher is its online engagement and new followers' acquisition rate. I then examine why firms that are more dissimilar online from their closest traditional competitors experience better online outcomes. To uncover the underlying mechanism, I analyze the content of all firm-tweets and categorize each tweet. The vast majority of firm-tweets fall into a hierarchy of 10 categories that correspond to social media affordances identified in prior research (Karahanna et al. 2018). Using 20,000 manually labelled tweets, I build two deep learning models that classify each tweet in the dataset into one of these 10 categories. The top-tier categories include tweets that emphasize cross-channel integration and tweets that involve users in co-creating value. The middle-tier categories include tweets that seek to create a community of like-minded users around specific series of events, product expert tips and product collections. The bottom-tier categories include tweets that seek to share content with users.

In examining the use of these different categories of tweets by firms, I find that firms that are dissimilar from their closest competitors are also better at leveraging Twitter's higher-level affordances for more value-added activities, as compared to firms that are more similar to their traditional competitors. These firms that are able to better leverage Twitter for higher value-added activities relative to their closest competitors also experience better outcomes online.

This study makes several important contributions. This study is among the first to examine why some firms outperform their close traditional rivals on SMPs. In doing so, this study contributes to growing body of research about firms' usage of social media, specifically how competitors use social media strategies. This study also contributes to research examining Isomorphism and divergence among firms in traditional context by extending this line of research to SMP competition. While most of the prior research on Isomorphism (DiMaggio and Powell 1983) focuses on similarities among competing firms, Beckert (2010) calls for research on divergent practices. This study responds to this call and seeks to understand the role of both similarity and divergence in rival firms' content strategies on social media platforms. This study also uncovers the mechanism related to how divergence in content strategies affects engagement as well as new followers' acquisition rate. While close traditional competitors experience competitive pressures to become more homogenic in their content strategies on SMPs, I show that firms that are able to differentiate themselves by better leveraging higher-level Twitter affordances relative to their rivals, experience better outcomes online.

In examining how firms differ in their use of Twitter, this study contributes to research literature related to SMP affordances. This study provides a hierarchy of Twitter affordances, wherein higher-level affordances that lead to more value creation have a significant and positive impact on online engagement.

The higher-level affordances allow using Twitter for co-creation and innovation by sourcing content and ideas from online users. While previous literature uses theoretical frameworks and case-based analyses to demonstrate such innovation (Mandviwalla and

Watson 2014, Mount and Garcia Martinez 2014), this study contributes to research literature on social media innovation by providing empirical support for the co-creation and innovation strategies described in the literature.

The findings of this study also have valuable managerial implications. My classification technique can be used to better understand firms' competitors on social media platforms and how the (dis)similarity in content strategies affects outcomes of interest. While SMPs offer the same set of affordances to all participants, not all firms are equally adept at leveraging these affordances. The findings show that it is not just the dissimilarity in firms' content strategies relative to their close competitors that gives firms an edge over their rivals, but their ability to leverage the higher-level affordances of Twitter that leads to better online outcomes for these firms. In identifying these levels of social media affordances, this study also provides valuable guidelines to managers seeking to leverage Twitter to design more effective strategies.

2.2 Related Research and Theoretical Background

Early research on social media platforms has focused mostly on individual users, users' networks and user-generated content. Users on social media start conversations, share content, build reputation, create relationships, and form online communities (Kietzmann et al. 2011, Dessart et al. 2015). Researchers have also analyzed the personality of users on social media platforms (Correa et al. 2010, Adamopoulos et al. 2015) and have linked certain personality traits to social media use and electronic word of mouth. Other studies have examined the connections among social media users and user networks on SMPs (Susarla et al. 2011, Zeng and Wei 2012). Finally, research focusing on users' behavior on SMPs has also studied user-generated content on social

media (Smith et al. 2012; Goh et al. 2013; Luca 2015). More recently, researchers have started to study firms' content strategies on online social media platforms. While this stream of research, as detailed below, examines specific usage of social media by firms, these studies do not focus on how firms' content strategies compare with those of their closest rivals and how their similarities or differences impact online outcomes.

This paper focuses on examining competing retailers' content strategies on Twitter and draws upon a well-established body of research that examines how firms compete in traditional settings. One of the major tenets of this stream of research is "Institutional Isomorphism" (DiMaggio and Powell 1983). Isomorphism has been defined as "a constraining process that forces one unit in a population to resemble other units that face the same set of environmental conditions" (Hawley 1968). Following the seminal work of Meyer and Rowan (1977), and DiMaggio and Powell (1983), a number of studies have focused on examining the antecedents of Institutional Isomorphism, including competitive pressures, which leads to homogenization, or *similarity* of institutions. Most of the work related to Isomorphism has focused on firms in traditional settings. There is some empirical support of the effect of competitive pressures on similarity of practices adopted by firms in traditional settings (Berrett and Slack 1999, Farndale and Paauwe 2007). More recent studies (for example, Beckert 2010) call for research that would focus on both Isomorphism and divergence. Beckert (2010) argues that the same set of mechanisms (i.e., mimetic, coercive and normative influences, as well as competition) could lead to either institutional homogenization or divergence depending on specific factors and circumstances. Beckert (2010) concludes that the theoretical challenge is to identify conditions under which these mechanisms

push institutional change towards similarity or divergence. This study adds to this stream of research by examining the similarity and divergence of traditional competitors' content strategies on online social media platforms.

Recently, researchers in Information Systems (Pant and Sheng 2015) have hypothesized that Isomorphism might also be observed in the Web footprints (firms' websites, online news, blogs, review platforms, co-searched firms and shared links among firms' websites) of competing firms. Pant and Sheng (2015) have found that the Web footprints of competitors are more similar than Web footprints of non-competing firms. This study adds to this emerging stream of work by examining the similarity of firms' content strategies and its implications for firms' outcomes on Twitter. Following on the lines of prior work that examines textual content shared by firms (Pan et al. 2015, Pant and Sheng 2015), I analyze the content shared by traditional competitors on Twitter to examine the degree of similarity in content among traditional competitors and further test how this degree of similarity affects the engagement of their followers on Twitter as well as a new followers' acquisition rate.

2.2.1 Social Media Content Categories

As noted earlier, recent research (Culnan et al. 2010, Swani et al. 2013, Einwiller and Steilen 2015, Lee et al. 2018, Bapna et al. 2019) on SMPs has begun to examine firms' use of such platforms. Firms post diverse content (news, updates about products or services, online promotions) on social media (Swani et al. 2013, Lee et al. 2018, Bapna et al. 2019), build brand communities (Culnan et al. 2010), handle complaints and manage their reputation (Einwiller and Steilen 2015). In examining the content posted

by firms on SMPs, researchers have identified specific categories of content shared by firms.

Brands, Products, and Firm-related Information: Malhotra et al. (2012), Malhotra et al. (2013), Pletikosa Cvijikj and Michahelles (2013), Luarn et al. (2015), Lee et al. (2018) and Bapna et al. (2019) find that firms use SMPs to share information about their brands and products (including product mentions, availability, or location), as well as firm-related information including firm milestones, partnerships, and achievements.

Events-related Information: Lovejoy and Saxton (2012), Malhotra et al. (2013) and Ashley and Tuten (2015) find that firms share information about offline and online events, including information about the theme, location and date of the events.

Discounts, Coupons, and Promotions: Malhotra et al. (2013), Lee et al. (2018), Bapna et al. (2019) find that firms use SMPs to share information about various offline and online deals, coupons, promotions and offers.

Expert Opinions and Product Tips: Parsons (2013) and Bapna et al. (2019) find that firms use SMPs to share expert opinions such as interviews, suggestions, and product tips, as well as advice and suggestions on how to use the firms' products.

Soliciting Questions and Opinions from Users: Malhotra et al. (2013) and Bapna et al. (2019) also find that SMPs are used by firms to solicit questions and opinions from their followers and fans.

Online-Offline Campaigns: Ashley and Tuten (2015) and Voorveld et al. (2018) find that some firms use SMPs to promote their advertising campaigns that originate offline (on TV, radio, newspapers, magazines or trade publications), while promoting their social media campaigns on traditional channels by using specific Web addresses or hashtags.

Interactive Contests and Soliciting User-Generated Content (UGC): Malhotra et al.

(2012), Parsons (2013), Ashley and Tuten (2015) and Gavilanes et al. (2018) find that firms also use SMPs to design and promote interactive contests and sweepstakes that encourage user participation. *Social Initiatives*: Lovejoy and Saxton (2012), Malhotra et al. (2013) and Lee et al. (2018) also find that firms use SMPs to share information about social initiatives performed by firms as well as calls for donations, volunteering, public service etc. Finally, Lovejoy and Saxton (2012), Malhotra et al. (2012) and Parsons (2013) find that firms also use SMPs to thank users and followers for engaging with the firms' brands and products, as well as to share holiday greetings.

Researchers (Rauschnabel et al. 2019) have also examined the role of social hashtags. Social tags facilitate content discovery and help users categorize, search, monitor, and participate in discussions based on user-defined tags. Nam et al. (2017) find that firms use social tags to monitor emerging trends, track engagement by topic and evaluate customers' views or sentiment towards a specific topic. From a brand perspective, firm-initiated tags serve a purpose of creating community of followers around specific products or campaigns (duel.tech 2018). Social tags boost interactivity by helping start and organize conversations. Specifically on Twitter, firms use hashtags to connect with other users that might search for specific topics or updates about events and to allow followers to join the conversation (Sundstrom and Levenshus 2017), as hashtags help users find like-minded community members. Social tags are also used by firms to collaborate with influential social media users as well as experts and bloggers (lonelybrand.com 2017) who can share helpful tips about firm's products. Finally, campaign hashtags are used for marketing campaigns or contests that run for a pre-defined time period (sproutsocial.com 2018). Firms leverage those hashtags

specifically for tracking contest entries and measuring conversation around a marketing campaign.

2.2.2 Social Media Affordances

Given that firms' successful use of content strategies depends on their ability to leverage the different capabilities of the social media platforms, this study also draws upon and builds on earlier work on social media affordances. As mentioned earlier, social media platforms are unique in a number of ways providing firms the ability to adopt strategies different from offline settings. Emerging research on IT affordances (see, for instance, Treem and Leonardi 2012, Yoo et al. 2012, Majchrzak et al. 2013, Volkoff and Strong 2013) in general, and social media affordances in particular, provides a good framework to examine the capabilities provided by SMPs. An affordance is defined as an action possibility that is available to an actor in the environment (Gibson 1986), and IT affordances are possibilities for a goal-oriented action afforded to specified groups of actors by technical objects (Pozzi et al. 2014). SMPs possess unique affordances that firms can exploit to creatively engage their followers (see Karahanna et al. 2018 for a review of social media affordances). While most of existing studies on social media affordances focus on organizational use of SMPs for internal communications and knowledge sharing, my work focuses on firms leveraging social media affordances for external audiences. Of particular relevance to my context are the affordances that are specific to firms using Twitter as a social media

platform – affordances such as *communication* and *self-presentation*, *meta-voicing*¹ and *relationship formation*, *interactivity*, *collaboration*, and *competition*.

As noted earlier, firms use Twitter to communicate with users and to share *content* by leveraging social media affordances such as *self-presentation* and *communication*. I term these lower-level Twitter affordances *content* affordances. However, social media affordances such as *relationship formation* and *meta-voicing* provide the capability for firms to not only share content online but also to create a *community* of users with similar interests around specific products or events through “continuous online communal knowledge conversations” (Majchrzak et al. 2013). On microblogging platforms like Twitter in particular, social tags (e.g., hashtags “#”) help to bring like-minded users around focal topics/products/events (sproutsocial.com 2018) and help users to create associations with other individuals or content (Treem and Leonardi 2012). I term these Twitter affordances *community* affordances.

Social media affordances such as *interactivity*, *collaboration*, and *competition*, enable firms to go beyond sharing content and creating a community of users to create value by engaging users in *co-creation* and combining multiple sources and channels. Researchers (Zittrain 2006, Autio et al. 2018) also refer to these higher-level affordances as *generativity*, or the ability to facilitate innovative inputs from large audiences, which enables organizations to engage customers in co-creative interactions, harness these innovative inputs for their offerings, and learn from these experiences. In addition to allowing individuals, groups, and organizations to co-create

¹ Meta-voicing is defined as “engaging in the ongoing online knowledge conversation by reacting online to others' presence, profiles, content and activities” (Majchrzak et al 2013).

services, applications, and content, generativity provides the ability to recombine data sources (Tilson et al. 2010) to produce new configurations and possibilities generating creative ideas that lead to innovation or produce overall value (Avital and Te'eni 2009). Prior research (Mandviwalla and Watson 2014, Mount and Garcia Martinez 2014) has used conceptual frameworks and case-based analyses to examine firms' use of SMPs for co-creation and innovation. In keeping with these studies, I term these higher-level Twitter affordances *co-creation* affordances.

While SMPs like Twitter enable a variety of affordances for all firms, these affordances are potential for action (Pozzi et al. 2014) and serve as possibilities for firms to achieve different objectives. However, these affordances need to be triggered and actualized (Strong et al. 2014) by firms to achieve desired outcomes. Not all firms might be equally adept at leveraging these social media affordances and consequently might differ in the outcomes achieved on these SMPs. In analyzing the content strategies adopted by firms on Twitter, this study also examines the firms' ability to leverage the different social media affordances and what these imply for firms' online outcomes. The literature on SMP affordances provides the framework for understanding what firms that are more similar to their closest offline competitors, do differently on Twitter and why they perform better on online outcomes relative to their competitors.

In addition to drawing upon prior work on social media affordances, this study also contributes to this stream of research by examining how close traditional competitors differ in their ability to leverage the different affordances provided by Twitter. This study also contributes to this stream of literature by showing how the differences in

leveraging social media affordances impact online outcomes of interest. The findings could be used to better understand how the effectiveness of firms' content strategies on SMPs are related to their ability to leverage specific social media affordances.

2.3 Research Context and Data

I begin by classifying traditional close competitors in the retail sector, as identified by Hoover's database (hoovers.com). I also verify if these firms are also close competitors from other sources such as Mergent (mergentonline.com) and Nasdaq (nasdaq.com)². In this study I focus on top three competitors, and these close competitors are also expected to have the highest offline similarity. I also confirm that these firms that are closest competitors have a high degree of similarity by analyzing their 10-K reports as detailed in the Section A2 of the Appendix.

To collect all tweets, I first open each firm's Twitter account Web page for each day, and collect each tweet ID, then using those tweet IDs I collect all tweets' text and metadata through Twitter API. Using Twitter API, I collect data for 199 retail companies from Russel 3000 list. The list contains the largest companies in terms of market capitalization. "Snowball" sampling was partially used to collect the data (Goodman 1961). For instance, I start with a focal firm A, and collect data on all top 3 competitors from Hoover's (B, C, and D); I then choose another focal firm from the (B, C, D) set and collect all top 3 competitors for each these focal firms – B, C and D. Hoover's database provides competitors unidirectionally, i.e., firm A might have top 3

² While the last two sources provide an average of 15 to 20 competitors for each focal firm, these sources do not identify the top competitors. However, Hoover's database specifically highlights top three competitors for each firm, and I confirm that these competitors are also listed as competitors in the other sources.

competitors B, C, and D, while firm B might have top 3 competitors C, D and E. If competitors' space is exhausted (i.e., no new competitors appear in the competitors set), I start with a new focal retail firm from the Russell 3000 list and repeat the process. I collect 2.42 million tweets (Table 2.1) for these retail firms for the period from January 2012 to August 2017. I focus my analysis on a single NAICS 2-digit sector, namely Retail Trade (NAICS sectors 44-45).

Out of 2.42 million tweets (Table 2.1), 893,525 tweets are firm-initiated. The rest (majority) are direct responses to customers' questions and complaints (which are marked by "@" tag at the beginning of a tweet), and retweets of tweets (minority) by a focal firm from other non-firm accounts (those tweets contain a tag of "RT @"). In the analysis I focus only on firm-initiated tweets, since those tweets should reflect firms' content strategies as well as the timing of those strategies with respect to their competitors³.

The Twitter data is a panel dataset with the tweets for the period from January 2012 up to August 2017. Each tweet has the text of the tweet, and such metadata as favorited tweets or "favorites" (a.k.a. likes), retweets, date/time, number of followers (a static number at the time of data collection, which is August 2017). I chose 2012 as a starting year for my analysis as I find that most firms in my dataset started actively posting

³ Few firms (e.g., Amazon, eBay) operate in multiple 2-digits NAICS sectors (at least the core sector is single for most firms). I retain those multiple-sector firms in my analysis and exclude them in the robustness checks to confirm that the results are consistent. Additionally, some firms have separate Twitter accounts for customer service (Q&A), job postings etc. In this paper I focus only on the primary official Twitter account for each firm that promotes all products across locations. For all firms in this study, I verify that the primary Twitter accounts are the ones with the largest total number of tweets and the number of followers. I explore the content of Q&A accounts, job posting accounts and other secondary Twitter accounts, and confirm that those accounts have very few tweets related to branding activities.

content on Twitter around 2012. Since this study focuses on analyzing the content shared by firms on Twitter, I ignore earlier time periods wherein the content shared by firms on Twitter is sparse.

Twitter API provides a static number of followers for each firm account in the tweets' metadata. Since I am interested in exploring drivers of a new followers' acquisition rate, I also collect followers' information for each firm account separately through a specific request of Twitter API. Next, for each follower ID I collect a follower profile. IDs and profile information are used to calculate the approximate date when each follower starts following a firm. These dates allow me to calculate the total number of followers and new followers by quarter as described in the Section A1 of the Appendix.

2.4 Methodology

2.4.1 Similarity/Dissimilarity

The methodological framework for calculating similarity/dissimilarity of content includes the following steps: first, I examine the pairwise similarity of traditional top competitors' content on Twitter. To this end, I first construct a term frequency-inverse document frequency (TFIDF) vector (Aizawa 2003) for tweets of each firm by quarter. Terms are unique unigram words of tweets from all firms in a quarter, also known as vocabulary/corpus. Frequency of occurrence of a term from vocabulary in each firm's tweets consists of the term frequency TF, and the number of a firm's tweets in which a term occurs determines the inverse document frequency IDF (as shown in the equation 2.1 below). Such numeric value reflects how important a word is to a document in a

collection/corpus. It can weigh down the effects of too frequent terms. Each firm is thus represented by a TFIDF vector.

Next, I use TFIDF vectors to calculate pairwise cosine similarity for each pair of firms for each quarter. I chose the cosine similarity measure, because it is one of the most commonly used methods for determining text similarity (Huang 2008), and because it addresses the problem of unequal corpus lengths. The cosine similarity is in the range from 0 to 1, where 1 is the most similar. Finally, I calculate average cosine similarity with top three competitors for each firm for each quarter.

To increase the quality of the data, I conduct several pre-processing steps, including stop words removal and non-ASCII character deletion as well as word stemming.

The formula for computing TFIDF value for a term t_i in document d_j is provided below (Aizawa 2003):

$$TFIDF(t_i, d_j) = tf(t_i, d_j) * idf(t_i) = \frac{f_{t_i}}{|d_j|} * \log \frac{N}{N_i} \quad (2.1)$$

where f_{t_i} represents the frequency of term t_i in document d_j ; $|d_j|$ is the number of words in document d_j . N_i – number of documents containing term t_i , and document d_j consists of all tweets of a focal firm in a quarter; N – total number of documents.

The cosine similarity calculation for two firms is formulated by the equation 2.2 (Pant and Sheng 2015):

$$sim(\vec{f}_A, \vec{f}_B) = \frac{\vec{f}_A \cdot \vec{f}_B}{\|\vec{f}_A\| \cdot \|\vec{f}_B\|} \quad (2.2)$$

where \vec{f}_A and \vec{f}_B represent TFIDF vectors for firms A and B. $\|\cdot\|$ is the length of a vector. The cosine similarity is a dot product of TFIDF vectors of two firms normalized by their lengths.

Cosine similarity metric has been extensively used in information retrieval literature (Huang 2008) and has gained popularity in other fields as well (Hoberg and Phillips 2016). Note that in this study I do not use Jaccard similarity because the tweet vector space is a continuous one and Jaccard similarity is specifically designed for a discrete space. Also, I do not use Euclidian distance due to its poor performance in high dimensional space.

Table A1 in the Appendix shows examples of cosine similarity scores for random tweets within a chosen category and across categories. Tweets' classification methodology is described below.

As can be seen from Table A1 in the Appendix, cosine similarity score is much higher within the same category of tweets (for example, within category “contests soliciting user-generated content”, see section 2.7) than across categories, suggesting that the cosine similarity algorithm correctly estimates topic similarity within content categories and topic differences across content categories. It is pertinent to note that the list of top traditional competitors may be asymmetric. For example, if a focal firm A has top three competitors B, C, and D on Hoover's, the focal firm B might have a slightly different set of top competitors C, D, and E. Thus, I consider average cosine similarity of each focal firm with top three competitors as defined by Hoover's database⁴.

⁴ In some of the cases where the official Twitter account for a third competitor does not exist or the account is very passive with very few tweets, I consider the average cosine similarity with the top two competitors. I also exclude 10 focal firms for which only one top competitor has an official Twitter account, and the other two competitors do not have Twitter accounts. Additionally, I perform the analysis for the subsets of focal firms for which all three top competitors have actively managed official Twitter accounts and achieve even stronger significant effect sizes.

2.4.2 Followers

Knowing the total number of followers that each firm has each quarter is important, since the higher number of followers is likely to produce more “favorites” and retweets. Thus, the dynamic number of followers could be used as a control variable (independent variable) in the econometric model. Additionally, using the number of followers by quarter allows calculating the number of new followers that a firm attracts every quarter. This variable is used as a dependent variable in later analyses.

Each tweet metadata provide only a static number of followers at the moment when metadata are requested from Twitter API. Bruns et al. (2014) provides a methodology to calculate the date when each follower starts following a focal firm. I obtain all followers’ IDs through Twitter API “GET followers/ids” command (developer.twitter.com). These followers’ IDs are returned in a very specific order - from the most recent to the earliest followers. I then collect profile information for each follower ID including username, screen name, description, location, and, most importantly, the date when a Twitter account was created by each user (i.e., follower). The ordering of followers and Twitter account date of creation are the two components used in the algorithm to calculate the “date of following” of a focal firm by each follower (see Section A1 in the Appendix).

I then create a count of the total number of followers in each quarter and the number of new followers by quarter. To do that, I count followers who started following a focal firm by the end of each quarter.⁵ Next, the number of new followers by quarter is

⁵ Alternatively, and more conservatively, one could calculate the total number of followers for each quarter using the beginning of each quarter as a cut-off. I use both approaches and obtain consistent results.

calculated as the number of followers at the end of the quarter minus the number of followers at the beginning of the quarter. I also validate this method by comparing the estimated number of followers by quarter with the actual number of followers (from a website web.archive.org) for a subsample of firms as described in the Section A1 of the Appendix.

2.5 Hypotheses and Econometric Specification

2.5.1 Hypotheses

Prior work on firm competitive strategies has found that competitive pressures lead to Isomorphism in competing firms' strategies (Berrett and Slack 1999), and such Isomorphism has been shown to positively impact firm performance (Brouthers et al. 2005). Particularly, when firms face competitive pressures under uncertainty, they might jump onto a bandwagon of adopting dominant strategies even if the outcomes of such strategies are ambiguous (Abrahamson and Rosenkopf 1993). For instance, when a firm enters a new domain, imitating an incumbent's "efficient" strategy could be beneficial for a firm's performance (Brouthers et al. 2005, Beckert 2010, Wu and Salomon 2016).

In the case of social media platforms, firms facing heightened uncertainties in decision-making with respect to social media content strategies, might jump onto a bandwagon of dominant strategies such as using similar online promotions, coupons and discounts that could drive online engagement and attract more new followers to a focal brand community. This strategy implies that competing firms already use very

efficient content strategies, and that the optimal choice for a firm would be to imitate those strategies. In this case, I would expect that:

Hypothesis 1 A. Isomorphism in firms' social media content strategies will have a positive effect on related outcomes.

A competing stream of studies has found that, in contrast to Isomorphism, firms that choose divergent strategies are likely to outperform their competitors (Badir et al. 2013). Institutional divergence is possible if institutional templates observed elsewhere are not considered legitimate solutions, thus, limiting the prevalence of imitation strategies (Beckert 2010). According to the resource-based view of the firm (Barney 1991), firms might choose to exploit resources that are valuable, rare, imperfectly imitable, and imperfectly substitutable (Farndale and Paauwe 2007) to differentiate themselves from their competitors and such differentiation could lead to better performance (Aulakh et al. 2000, Badir et al. 2013). Further, competition may support the development of diverse organizational models that allow for niches in which companies can specialize (Beckert 2010) because firms are equipped to do different things and can therefore capitalize on “their indigenous sources of strength” (Guillén 2003).

Thus, in the context of social media platforms, firms might forgo imitation of potentially inefficient dominant strategies of competitors and strategically choose (Farndale and Paauwe 2007) competitive divergence, and leverage unique Twitter affordances to experiment with Twitter content and use Twitter to differentiate themselves from top competitors, and such divergence will be more engaging for online users and will attract more new followers. If so, I would expect that:

Hypothesis 1 B. Divergence in firms' social media content strategies will have a positive effect on related outcomes.

Thus, whether Isomorphism or divergence in social media content strategies leads to better online performance, remains an empirical question – one that I examine in this study.

2.5.2 Econometric Specification

To test the effect of similarity on online engagement, I estimate the following fixed effects model:

$$Y_{it} = \beta_0 + \beta_1 \text{Similarity}_{it} + \beta_2 \text{Tweets}_{it} + \beta_3 \text{Followers}_{it} + \alpha_i + \delta_y + \delta_q + u_{it} \quad (2.3)$$

The main independent variable (similarity) in the specification is the average cosine similarity with top 3 competitors for each firm for each quarter. I choose the quarter as the period in the panel data, since I believe that quarterly data will have enough tweets to measure similarity in content strategies for each pair of firms even if posting frequencies for firms differ (for example, some firms post once per day, some firms post less frequently). The outcome variable (Y_{it}) represents engagement on Twitter, as measured by the total number of “favorites” and total number of retweets per quarter. I report the results separately for “favorites” and retweets. Another outcome variable (Y_{it}) is the number of new followers gained in a focal quarter. The number of tweets (Tweets), and the dynamic number of followers (Followers) for a given firm for a given quarter are control variables. The variable α_i represents firm fixed effects, the variable δ_y represents year fixed effects, and the variable δ_q represents quarter fixed effects.

I leverage the panel structure of my dataset to include the firm fixed effects in the model. Firm fixed effects control for time-invariant unobserved heterogeneity across

firms. Additionally, I add year and quarter fixed effects to the model to control for potential time trends in firms' content strategies.

Table 2.2 shows descriptive statistics of the variables in the models⁶.

As a robustness check, I add additional controls to the model such as images and videos. Out of all 893,525 firm-initiated tweets, 357,926 tweets have an image as part of the tweet, and 11,663 tweets have a video as part of the tweet. To add control variables, I calculate number of images and videos for each firm for each quarter. While control variables have positive effect on engagement, they do not affect the main results of the model related to similarity/dissimilarity. Also, it is pertinent to note that adding such control as tenure for each firm (number of days since opening a Twitter account) in each quarter does not affect the main results.

An additional robustness check involves adding Google Search Trend scores (trends.google.com) for each firm for each quarter to control for potential unobserved offline marketing campaigns (Ghose et al. 2012, Wang et al. 2018) that could affect online engagement on Twitter. Adding these Google Search Trends to the model does not affect the main results.

Additionally, I want to make sure that similarity or divergence in content of a focal firm is driven by the focal firm and not driven by competitors' content strategies. Thus, as a robustness check, I restrict the sample only to observations (firm-quarters) when competitors have very little changes in their content strategies as compared to the

⁶ It should be noted that correlations among independent variables do not go beyond 0.33 (the range is from -0.33 to 0.25) in absolute values. Additionally, Variance Inflation Factors (VIF) for all variables are lower than 2 (which is much lower than the "problematic" value of 10, and lower than a more conservative "problematic" value of 4). Thus, multicollinearity is not an issue in these specifications.

immediate prior quarter. I can control for competing firms' content strategies after performing classification of all tweets into 10 main categories of content as described in Section 2.7 below. In this robustness check I find the results to be highly consistent with the main findings.

I also perform additional robustness checks where I compare the online similarity social media metric with the offline similarity metric calculated using one of the annual reports issued by SEC, 10-K (a comprehensive summary of a firm's financial performance). This analysis helps reinforce my identification of traditional competitors. While the main analysis identifies traditional competitors based on Hoover's classification, analyzing the content of firms' 10-K reports serves as a useful robustness check to verify if these traditional competitors are indeed similar offline (Hoberg and Phillips 2016). I find that top competitors have higher 10-K similarity than other (distant) competitors, and that other competitors have higher 10-K similarity than non-competitors. Thus, my identification of top competitors is supported by the 10-K analysis. The details of this analysis are provided in the Section A2 of the Appendix.

To address any endogeneity concerns related to similarity, I employ a system GMM model using Arellano–Bover/Blundell–Bond linear dynamic panel-data estimation with lagged dependent variables as instruments. I use a Stata xtdpdsys 2-step estimation with 2 lags of the dependent variable and with robust standard errors suggested by Windmeijer (2005). One of the prerequisites of using dependent variable lags as instruments is that there should be no second-order autocorrelation of residuals. I test that condition and confirm that the autocorrelation of the second order is not present in all model specifications (the p-value higher than 0.14). With robust standard errors, the

Sargan test of overidentifying restrictions is not calculated. But this test is conducted with “gmm” errors (Stata command “estat sargan”) and a 1-step estimation, and that test is passed in all model specifications ($\text{Prob} > \chi^2 = 0.99$). It is pertinent to note that adding lags of independent variables (for the number of tweets, dynamic number of followers and similarity) does not change the results. Further, to address potential reverse causality issues (when online engagement can affect differentiation strategies), I use a cross-lagged panel model with fixed effects with maximum likelihood estimation (xtdpdml package in Stata, see Allison et al. 2017) that has been shown to address reverse causality issues better than other models (Leszczensky and Wolbring 2019). I find the results to be consistent with the GMM model.

Additionally, I use a 2-stage least squares estimation with heteroskedasticity-robust standard errors with the following instrumental variable – moving average similarity with top competitors in the last 2 to 5 quarters before the current quarter (excluding the current quarter). Using 2, 3, 4 or 5 prior quarters gives consistent results. The instrumental variable (IV) should affect a firm’s propensity to be dissimilar in each current quarter and should not directly affect relevant outcomes (online engagement and new followers’ acquisition rate) in each current quarter. The only effect of IV on outcomes should come from dissimilarity. I believe that this instrumental variable will satisfy exclusion restriction criteria, since dissimilarity in prior quarters will most likely not have a direct effect on online engagement in the current quarter since most firms post a large number of tweets per quarter (553 tweets per quarter, on average, see Table 2.1). To see tweets in prior quarters and to “engage” with them, a user has to scroll down more than 50 pages, which is time-consuming. I also exclude the immediate prior

quarter and find the results to be consistent. Thus, dissimilarity in prior 2 to 5 quarters will likely affect dissimilarity in a current quarter but will not directly affect online engagement in a current quarter. The first stage of the 2-stage model is highly significant. Additionally, the null hypotheses of under-identification and weak identification (using Kleibergen-Paap rk LM statistic) are rejected, and Hansen J statistic cannot be used for cases of 1 IV for 1 endogenous regressor.

2.6 Results

Figure 2.1 provides an example illustrating the degree of similarity among traditional competitors in their content shared on Twitter. Consider the 4 top competitors (as per Hoover's database): focal firm A and its top 3 competitors B, C, and D.

These 4 firms represent a single offline cluster (blue rectangle in Figure 2.1) and have a high degree of similarity with each other based on their 10-K reports as well (10-K similarity is not shown on Figure 2.1). On examining their pairwise cosine similarity of content on Twitter, I see that firms B, C, and D are similar to each other on Twitter based on the content analysis. However, firm A is much more distant from (i.e., dissimilar to) the 3 other firms.

I find that such a pattern – that Twitter content of some firms is more similar to the content of their traditional competitors, while Twitter content of some other firms is less similar (or dissimilar) to the content of their traditional competitors - is observed in a large number of subsets of top competitors in my dataset. For example, among the 199 firms in the Retail Trade sector one or two dissimilar firms (outliers) were clearly seen in more than 50% of the subsets of competitors. Figure 2.2 demonstrates this pattern for 2 subsets of 4 top competitors where the outliers are the green nodes.

I then examine how the degree of similarity of a focal firm's content with content of its closest competitors is linked to the outcome variables, namely, online engagement and the acquisition rate of new followers. To recall, the independent variable is the average cosine similarity of content with top three competitors for each firm for each quarter, and the outcome variables are the total number of "favorites" for each quarter, the total number of retweets for each quarter and the number of new followers for each quarter.

The results (Table 2.3) show that the more distant a firm's content is from its traditional competitors (i.e., the lower the degree of cosine similarity of content with its traditional competitors), the higher is the online engagement on Twitter.

The specification with fixed effects for each firm ("within-variation") has a negative coefficient for similarity, which indicates that higher similarity leads to lower online engagement. The GMM and 2SLS models support the results. In other words, the more *dissimilar* a firm is from its closest traditional competitors, the higher is its online engagement⁷. I find the same effect for retweets as for "favorites". In examining the effect size, for example, the value for "favorites" in the fixed effects model is "-31,475" (Table 2.3), which suggests that, given the average value is 6,375 (Table 2.2), increasing the similarity of content with top competitors from 0 (no similarity) to 1 (perfect similarity) would decrease the number of favorites by 31,475. Table 2.2 also

⁷ It is important to note that the number of focal firms in the Retail Trade sector is 156 (Table 2.3), whereas the total number of competitors in that sector is 199 (Table 2.1). Five firms in my sample have millions of followers that could not be processed by the algorithm to calculate the dynamic number of followers. I drop these firms from the focal firms list (but they still appear as competitors for other focal firms). Additionally, the sampling is performed partially by a "snowball" method, and I focus on firms that have top competitors within the retail sector resulting in 156 focal firms out of a total of 199 firms.

shows that the average cosine similarity of a firm's content with content of its top competitors is 0.09. Thus, it is more realistic to expect that when dissimilarity increases by 0.1, the number of "favorites" (per quarter) is expected to increase by 3,147.5 (or 49.3% from the average value), which is both statistically and economically significant.

When it comes to new followers, the GMM model with 2 lags of the dependent variable as instruments has a significant (a p -value slightly higher than 0.05) negative coefficient of similarity, which indicates that firms with more dissimilar content relative to their closest traditional competitors attract more followers online. The effect in the 2SLS model is negative and significant with the p -value < 0.05 . Additionally, the control variables (dynamic followers and tweets) have positive coefficients in all model specifications, as expected. The exception is non-significant coefficients for the dynamic number of total followers affecting the new followers' acquisition rate in the fixed effects and GMM models.

The number of observations in the GMM model is 3,197, due to the dropping of 2 quarters of data, since 2 lags of dependent variables are used. The 2SLS model uses an instrument (moving average similarity in previous 2 to 5 quarters), so one quarter is dropped for all 156 firms (3,353 observations remain).

It is also pertinent to note that, for example, about 600 tweets (out of 893,525 in my dataset) generate more than 10,000 "favorites" and/or retweets. Those "viral" tweets may be responsible for a part of the identified effect. As a robustness check, I re-estimate the models after removing top 1-5% of quarters with the highest scores for "favorites" and/or retweets and obtain consistent results.

I also perform a series of robustness checks that involve restricting the sample to firms with few followers, restricting the sample to firms with many followers, removing firms with millions of followers or very few followers, as well as splitting samples into more actively posting firms (300 and more tweets per quarter) and less actively posting firms (less than 300 tweets per quarter). I choose those splits based on distribution of relevant variables. In all these cases I find that the more dissimilar a firm's tweets are from its closest competitors, the higher is its online engagement. A significant effect of dissimilarity on attracting new followers is supported by the GMM model and the 2SLS model.

Overall, the positive correlation of “favorites” with retweets is 0.89, so those two outcome variables mirror each other. More importantly, I find that higher dissimilarity attracts more new followers.

Another robustness check relates to whether Twitter changed its timeline algorithm to feed tweets to users in a particular way (for example, using a feature “show me the best tweets first” etc.). If tweets are shown in some nonrandom manner, it could be the case that most dissimilar tweets might somehow have more impressions than similar tweets. Twitter changed its timeline algorithm in March 2016 (socialmediatoday.com 2016). Prior to that, tweets were shown in reverse chronological order (starting with the most recent tweets). As a robustness check, I restrict the data to years 2012, 2013, 2014 and 2015, and the results are consistent. Thus, I believe that the Twitter feed algorithm does not impact my results.

2.7 Why Do Dissimilar Firms Perform Better?

My results, thus far, show that the more dissimilar a firm is to its closest traditional rivals with respect to its Twitter content, the better are its outcomes on Twitter. However, it is not clear what is specific about the dissimilarity in content that leads to better outcomes. To understand the drivers of why dissimilarity in content strategies is associated with better outcomes, I examine what dissimilar firms do differently compared to their close rivals. As described earlier, prior research identifies 12 categories of content that firms share on social media platforms. I analyze my dataset of firm-initiated tweets to confirm these categories of content identified in prior work. Content classification is conducted in two stages in keeping with prior research (Ashley and Tuten 2015, Bapna et al. 2019). In the first stage I use 20,000 random tweets to perform manual classification of firm tweets. Next, I rank categories in terms of frequency, and I choose 10 most frequent categories⁸ that represent more than 75% of all tweet content. I use three expert raters to validate the identified 10 categories. The inter-rater reliability (percentage of agreement) among experts for the manual classification step for 10 categories is 87%. Content that is not a part of the 10 most frequent categories is assigned to a “Misc.” category.

The second stage (similar to the approach of Lee et al. 2018) involves building deep learning models using 20,000 labelled set as a training sample and using those models to classify all tweets in the dataset. Tweets with low predicted probability of belonging

⁸ While the total number of tweet categories is larger, the 10 categories used in the analysis are the most frequent. The remaining tweets (about 25% of all tweets) are categorized under “Misc.”. These mostly include tweets containing “thank you” messages, “birthday greetings”, information about social initiatives, and public service announcements.

to one the 10 most frequent categories are labelled with a “Misc.” category. More details about classification and deep learning models are provided below and in Section A3 of the Appendix.

Content: The five most common categories consist of firms sharing content relating to brand/*product information, product usage tips, questions, events, and coupons and promotions* through their tweets. These tweets that are primarily focused on firms sharing content with users leverage Twitter’s affordances such as self-presentation and communication.

Community: The next set of categories include tweets where firms seek to create a community of like-minded users (sproutsocial.com 2018). All of these tweets involve the use of Twitter hashtags “#”. The categories include firm tweets relating to “*#expert tips*”, “*#product collections*”, and “*#special events*”. The use of Twitter hashtags “#” enables firms to create a community of like-minded followers and foster interactivity around a focal campaign/contest/event/product (sproutsocial.com 2018). In contrast to tweets without a hashtag that are primarily used to disseminate information about products/events/promotions, the use of hashtags serves as a mechanism for enabling the realization of higher-level affordances, such as relationship formation, and meta-voicing.

Co-creation: The final set of tweet categories include tweets wherein firms seek to involve users more actively to help create and share content relating to their offerings. Twitter hashtags are used to help users not only create a community around specific topics but also involve them using specific campaigns and contests. These categories include “*#offline-online campaigns*” and “*#contests soliciting user generated content*”

that leverage Twitter affordances including interactivity, collaboration, and competition.

A detailed description of these categories with examples are provided in Table A2 in the Appendix.

Next, using these labeled tweets, I train two deep Recurrent Neural Network (RNN) models with Long Short-Term Memory (LSTM) using the Keras package (keras.io) of python – one model for 5 categories with hashtags (accuracy is 86% on the 20%-hold-out test set, average precision is 83.7, average recall is 87.9) and another model for 5 categories without hashtags (accuracy is 87% on the 20%-hold-out test set, average precision is 84.9, average recall is 88.8). The details of the deep learning models are provided in the Section A3 of the Appendix. The two models are used to classify all 873,525 firm-initiated tweets, which are 893,525 tweets minus the training sample of 20,000 tweets. I separate all 873,525 tweets into 2 groups (see Table 2.5 for distribution of tweets in two groups) – tweets with hashtags and tweets without hashtags and use two trained models to classify tweets into related categories. All later analyses are performed on all 893,525 firm-initiated tweets. Tweets with low probability (lower than 40%) of belonging to any group are assigned to the category “Misc.”. Thus, the category “Misc.” includes tweets with hashtags and without hashtags, which were not assigned to any of the 10 categories. As a robustness check, I adjust the cut-off for the “Misc.” category from 20% to 30% to 40% and to 50% and the results are consistent.

Hypotheses. Figure 2.3 illustrates the classification of the 10 tweet categories into a hierarchy of three tiers. The bottom tier consists of firm-tweet categories where the focus is on communicating information and sharing *content* with all users. All firms in

my sample use Twitter affordances including *self-presentation* and *communication* to share content with their users. As mentioned earlier, I term the tweets' categories that leverage this basic set of affordances "*content*" categories.

The categories in the middle-tier leverage *relationship formation* and *meta-voicing* affordances and enable firms to not only share content but also create a *community* of users around focal themes ("*#Expert tips*", "*#Product collections*", "*#Special events*"). I term the tweets' categories that leverage these mid-tier affordances "*community*" categories.

The categories in the top tier involve *cross-channel integration* ("*#Offline-online campaigns*") as well as engaging users in *co-creating value* ("*#Contests soliciting UGC*"). The affordances that correspond to that tier are interactivity, collaboration, and competition. Firms leverage Twitter's interactivity affordance to create offline-online campaigns when users are encouraged to share their ideas. Firms leverage collaboration and competition affordances to design contests soliciting users to upload user-generated content. I term the tweets' categories that leverage these top-tier affordances "*co-creation*" categories.

The hierarchy of tweet categories illustrated in Figure 2.3 captures the fact that the higher-level categories include affordances that build on affordances in lower levels. While self-presentation and communication affordances constitute the most basic tier of affordances, to create a community requires a firm to leverage affordances in the "*content*" layer and use a hashtag frequently to focus users' interest on a specific topic (wikihow.com 2019). Likewise, to initiate a contest soliciting UGC or an offline-online interactive campaign, a firm needs to use "*content*" affordances as well as create a

“community” around the content by introducing a hashtag and posting it frequently and encouraging users to share ideas and other UGC using this specific hashtag (shortstack.com 2018).

Further, higher-level affordances are more demanding for firms in terms of effort and investments. While the basic “content” affordances require some investments into creating content and communicating with users, these are usually similar to the content that is promoted by the firm in other non-social media channels. In contrast, leveraging “community” affordances requires not only carefully selecting the content to create a community around, but also additional investments in introducing and promoting a hashtag for specific content across multiple user touch points (blog.hubspot.com 2019). Leveraging “co-creation” affordances requires further coordination of “content”, “community” and “co-creation” strategies, and requires even more time and resources (Roberts and Piller 2016). In addition to requiring higher investments from the firm, higher-level affordances are also more difficult to execute successfully.⁹

Given the higher cost and resources required by higher-level affordances, I expect fewer firms to leverage the higher-level community and co-creation affordances as compared to firms leveraging the lower-level content affordances. Thus, I hypothesize that *firms that leverage higher-level affordances are likely to be more dissimilar from firms that leverage only the affordances at lower levels in the hierarchy*. Further, since the higher-level affordances are aimed at creating a community and getting users more

⁹ For example, in 2014 Puma invited Twitter users to tweet the hashtag #ForeverFaster, following which an automatic “personalized message” would be generated signed by one of Puma's celebrity brand ambassadors. However, Twitter users discovered that they could change those personalized messages and made it seem as if superstars such as Usain Bolt, Radamel Falcao and Cesc Fabregas had written offensive messages (dailymail.co.uk 2014).

involved with the firms' content, I hypothesize *that higher-level affordances will have higher online engagement and will attract more new followers.*

As highlighted by previous studies, the main opportunities for innovation on social media come from creating communities (Culnan et al. 2010) of dedicated online users and involving those online users in value co-creation through crowdsourcing of ideas (Mandviwalla and Watson 2014) and user-generated content (Roberts and Piller 2016). As stated by Roberts and Piller (2016), it is not enough to just set up a social media profile and wait for users to come, engage and co-create. Firms need to grow community of followers, and then use data mining methods to determine product trends and customer sentiment. The next level of engagement would imply running ideation contests and rewarding top users. Next, a firm might invite users to participate in the design of a product. Finally, once a product is launched, a firm might use voices of active users to promote positive online word-of-mouth across the community of followers. This is also consistent with Helms et al. (2012), where the authors describe five active social media innovation strategies that are ordered according to user participation levels: *general community engagement, ideas competitions, interactive value creation, participatory design, and product design*. By creating a community, a firm stimulates online users to share experiences with like-minded people. In the context of Twitter, firms have a big community of followers, and use hashtags to create sub-communities of users with closely related interests that form around an event, a product or a campaign. Next, firms can create contests on Twitter to solicit user-generated content. These contests are designed to solicit novel creative ideas in the form of general design suggestions or desired product features. An interactive value

creation occurs when, for example, a firm runs a competition for the best post, best photo caption or best photo with its products. Such user-generated content receives votes from other users and might be used by a firm later in its marketing materials online and offline. Participatory design and product design involve even more focused crowdsourcing campaigns with the goal of soliciting new ideas that could be used to launch new products.

Thus, based on the above, I hypothesize that *“content” affordances will have the lowest level of user engagement, because they do not actively involve users in innovation*. The “community” affordances are expected to be related to a *general community engagement* strategy, and, thus, will have higher online engagement than “content” affordances. Finally, “co-creation” affordances that include *ideas contests, interactive value creation, participatory design and product design* strategies are expected to have the highest engagement among all levels.

To test the hypotheses, I use cosine (dis)similarity scores for 156 firms for each quarter and calculate proportion of tweets in each category for each firm. I estimate a beta regression model with a logit link, where the dependent variable is the similarity score, and the independent variables are a percentage of tweets in each category (all 10 categories comprise 100%). Table 2.4 shows the results - higher dissimilarity (i.e., negative coefficients for similarity) is associated with an increased usage of higher-level affordances.

Table 2.5 shows a comparison of normalized engagement for all categories. I find that tweets in the categories in the higher tiers have higher engagement than those in the lower tiers.

Table 2.5 has all 10 categories listed in the order of corresponding tiers (1-3). I find that for “favorites”, categories in tier 1 (top) have higher normalized engagement than categories in tier 2, which in turn have higher values than those in tier 3 (bottom). Regarding retweets, the general pattern is the same as for “favorites”, with one exception. Tweets in the category “coupons and promotions” in tier 3 have a higher number of normalized retweets than the two categories in the 2nd tier (#expert tips and #special events).

The category “#expert tips” in Table 2.5 has a higher number of tweets compared to other higher-level categories. This is potentially due to the higher number of fashion firms in my sample of 199 retail firms. Based on the observation of the dataset, fashion firms use style tips extensively. Nevertheless, as shown in Table 2.4, dissimilar firms are found to use a higher proportion of “#expert tips” compared to similar firms.

I have hypothesized that the firms leveraging higher-level affordances are likely to be more dissimilar from firms that leverage only the lower-level affordances. My results support that hypothesis.

The second hypothesis was related to higher engagement and new followers’ acquisition rate for higher-level affordances compared to lower-level affordances. Table 2.5 shows that higher-level affordances have higher normalized engagement. To formally test the effect of higher-level affordances on related outcomes, I use the earlier instrumental variable (moving average similarity in previous 2 to 5 quarters) to instrument for the usage of higher-level categories and test whether such usage affects online engagement and new followers’ acquisition rate. The instrument passes validity checks for a good instrument. Specifically, the first stage of the model is highly

significant. Additionally, the null hypotheses of under-identification and weak identification (using Kleibergen-Paap rk LM statistic) are rejected. Table 2.6 shows the results.

According to Table 2.6, higher-level categories have positive effect on online engagement and new followers' acquisition rate. Regarding the lower-level categories, the instrument (moving average similarity with top competitors in previous 2 to 5 quarters) cannot be used for the usage of those lower-level categories since similarity is not associated consistently with the usage of lower-level categories (see Note to Table 2.4). Thus, the results for the GMM model (not reported in Table 2.6) show that higher proportion of tweets in those lower-level categories either affects relevant outcomes negatively or does not have an effect (except for "coupons and promotions" that seem to attract new followers).

Additionally, I perform a mediation analysis (see Section A4 in the Appendix) and show that dissimilarity affects online engagement and new followers' acquisition rate through the usage of higher-level affordances (full mediation). The mediation analysis confirms my hypotheses.

To summarize, I examine the mechanism behind why dissimilarity in content strategies positively affect firms' online outcomes on Twitter. To this end, I classify all tweets into 10 categories that are mapped onto social media affordances identified in prior research. Higher-level categories with higher-level affordances require more resources and coordination, making it costly for other firms to successfully implement them. I expect that fewer firms will use higher-level affordances as compared to lower-level affordances that are easier to leverage and coordinate. The analysis of tweet

categories supported by the mediation analysis confirms that dissimilar firms differentiate themselves from top competitors by adopting higher-level content categories with higher-level affordances. Further, I find that higher-level affordances have higher online engagement and new followers' acquisition rate compared to lower-level affordances.

2.8 Implications and Conclusion

This study seeks to understand the similarity/dissimilarity of content strategies in online social media platforms of close traditional competitors and consequences of such content similarity/dissimilarity for outcomes in these platforms. In doing so, my study introduces a new measure of online social media competition based on similarity of content with top competitors. I show that both Isomorphism and divergence could be observed in online social media content strategies. I find that divergence in content strategies from a firm's closest rivals leads to higher online engagement and attracts more new followers for the focal firm. While earlier research (Pant and Sheng 2015) shows that close traditional competitors are more likely to adopt similar content strategies compared to other firms, this study focuses on the differences in content strategies among close traditional rivals. I find that although close traditional competitors have a high degree of similarity offline, there is greater dissimilarity in their online content strategies on Twitter, and that these differences have important consequences for firms' Twitter outcomes.

To understand the underlying mechanism behind the positive effect of divergence on related outcomes, I classify content on social media into 10 categories that map onto 3 tiers of social media affordances. I find that dissimilar firms not only adopt different

content strategies on Twitter as compared to their closest rivals but are also adept at leveraging higher-level social media affordances and that higher-level social media affordances lead to higher online engagement.

This study makes several important contributions. First, there is a growing body of research literature related to how firms use social media platforms. However, there is very little research related to competitors' content strategies on SMPs. My paper seeks to fill the gap by exploring traditional competitors' strategies on Twitter. Next, the results of this study make contribution to research literature on Isomorphism of traditional competitors, and specifically to research related to dynamic modern methods of competitors analysis using Web footprints of rivals. The imitator (Wu and Salomon 2016) or innovator (Zheng Zhou 2006) dilemma has been described in the research literature related to traditional firms in offline channels. In the context of Twitter, firms, facing competitive pressures, might experience uncertainty in decision-making about SMP strategies and imitate their competitors' dominant content strategies. Alternatively, firms have an option to strategically choose to be divergent in content strategies and use Twitter innovatively. I find that it is divergence that leads to better online outcomes. Thus, traditional rivals can overcome the pressure to be isomorphic on social media platforms and use the strategy of divergence that leads to higher engagement with their followers and attracts more new followers.

My findings related to how divergence leads to related outcomes make contribution to research literature on the use of social media affordances by rival firms. While previous research identifies social media affordances, it does not explore the differential impact of affordances on outcomes. I propose a 3-tier framework of tweet

categories and affordances and show that higher-level categories with higher-level affordances not only include lower-level categories with lower-level affordances, but also create additional value for the dissimilar firms that use higher-level categories in higher proportion compared to their more similar rivals. Hierarchically, the value of content categories increases when firms shift their content strategies from sharing content to creating community of like-minded users to co-creating value with online users.

I find that the most engaging types of tweets are related to co-creating and innovating by leveraging higher-level affordances of Twitter. Earlier research on social media innovation has used theoretical frameworks and case-based studies to demonstrate the value of co-creation and innovation on social media platforms. Note that the first step to start the co-creation process on social media platforms is to create a community of like-minded users (Culnan et al. 2010). Helms et al. (2012) propose different levels of user involvement in social media innovation – ranging from community engagement to participatory design and product design. In this study I provide empirical support for these propositions by showing that dissimilar firms use Twitter innovatively, i.e., create communities around specific hashtags and source ideas and content from online users and that such co-creation strategies outperform other types of social media strategies. This study adds to this growing stream of research on social media innovation by showing that firms that leverage higher-level affordances differentiate themselves from their closest competitors and reap the benefits in the form of higher online engagement and higher number of new followers.

This study also contributes to the research literature on early mover advantages. While the close traditional competitors in my sample are well-established firms that compete closely with their traditional rivals, I find that not all of them are equally adept at leveraging the different affordances of Twitter. I find that early movers, who have been able to leverage Twitter's higher-level affordances, better than their closest competitors, experience better online outcomes. Whether these early mover advantages are sustainable would be an interesting topic for future research.

My findings have some important managerial implications. Managers can exploit the ranking of competitors by dissimilarity of content on social media platforms to determine which of their top traditional competitors have a potential social media competitive advantage. The dissimilar rivals are more likely to leverage higher-level social media affordances and experience better online outcomes. Additionally, managers can use my hierarchy of content categories and affordances to better design their social media strategies. While bottom-tier affordances are relatively less-costly for firms, their impact on engagement is limited. Higher-level affordances require higher investments from firms, but those investments pay off in the form of higher online engagement and higher number of new followers. Based on my findings, the most valuable content strategies involve users in co-creation and innovation activities that are not only highly engaging for followers of brand communities but could also be used for growing brand communities, and eventually for new content and new products development. But, importantly, the most effective co-creation and innovation strategies should be built on top of other strategies such as community building strategies around products, events and experts. Prior studies have shown that social media engagement

positively affects firms' brand value and sales (see, for example, Kumar et al. 2013). Thus, it is important for managers to know capabilities of social media and how they can leverage those capabilities for increasing online engagement. Another implication is that firms not only need to leverage the interactive affordances of SMPs, but also need to provide mechanisms for users to keep track of, and engage with, the firms' interactive online campaigns over time.

This study is not without limitations. First, I explore firms' social media content strategies in the retail sector. Future research could investigate firms' strategies in other industry sectors such as Arts, Entertainment, and Recreation or Finance and Insurance. Firms' social media strategies in these sectors might differ from strategies in the retail sector, and it would be interesting to examine how firms in these sectors compete on SMPs.

Next, in this study I do not investigate details of followers' behavior for each firm's official Twitter account. Future research could explore each firm's followers' activity to see if dissimilar firms not only better engage their current followers but also attract more engaged distinct loyal followers (brand fans) that follow only a focal dissimilar firm and do not follow competitors. Those loyal brand followers might be a source of strong online and offline word-of-mouth and might act like brand ambassadors. Future research can further explore the detailed role of social media affordances in creating and sustaining higher online engagement, as well as the value to firms of combining online strategies with offline marketing campaigns.

Additionally, future research may further explore innovation strategies on social media platforms. More empirical work is needed to relate various innovation strategies to unique social media affordances.

Future extensions to this work can examine whether online strategies impact offline metrics, for example, firm sales or stock prices. Finally, this study focuses on firms' use of one social media platform, Twitter. Future work can analyze firms' strategies on other social media platforms, such as Facebook, and examine if firms' competitive behaviors are similar across these platforms.

Chapter 3: Synthesizing Winning Strategies: What Differentiates Experienced Designers in Crowdsourcing Markets?

3.1 Introduction

Online crowdsourcing platforms have grown in popularity in recent years with a plethora of platforms for a variety of tasks including crowdfunding, open innovation platforms, and crowdsourcing platforms for design, among others. In keeping with the growth of these online marketplaces and platforms, there has been an increasing interest among academicians as well as practitioners in understanding various aspects of these platforms. Given that several online crowdsourcing platforms, especially those for innovation and design tasks, use contests to solicit submissions from designers as well as to decide on a winner, prior research has examined various aspects of the design of such contests and their impact on outcomes for different participants including the contest holder as well as the designers.

Prior research on designers' behaviors in online crowdsourcing platforms focuses on factors such as timing of entry of designers (Yang et al. 2010), the impact of prizes (Araújo 2013), competition (Shao et al. 2012) and feedback (Wooten and Ulrich 2017) on outcomes of interest including their likelihood of winning a contest. A few studies also examine factors that differentiate successful entries from others. In particular, some of these studies examine the behaviors of experienced designers and how they differ from the behaviors of less-experienced designers. These studies (Yang et al. 2010, Khasraghi and Aghaie 2014) find that experienced designers are more likely to make submissions earlier or later in the contest and are more likely to win as compared to less-experienced designers. However, what is it that experienced designers do that leads them to be more successful is less well understood in these markets.

Understanding what strategies work best and why, is crucial for both academicians and practitioners.

This study focuses on open contests on a crowdsourcing platform for design tasks where the “open” nature of contests (i.e., visibility of prior design submissions) and public feedback (i.e., star ratings provided by clients) provide an opportunity for designers to leverage information from prior design submissions made by other designers. This study is motivated by and builds on two theoretical foundations. The first stream of research focuses on information spillovers. To the best of my knowledge, very few papers have explored information spillovers in open crowdsourcing contests for design tasks, potentially due to difficulty of measuring spillovers directly. The second stream of research is related to the possibility of recombining information spillovers from prior design submissions (LaToza et al. 2015). Prior research (Fleming 2001, Brynjolfsson and McAfee 2014) on “recombinant innovations” finds that a majority of inventions happen by recombining prior inventions. Fleming (2001) and LaToza et al. (2015) find that inventors and designers, as they gain experience, learn which components they need to recombine and what the optimal recombination strategies are. In keeping with this, I hypothesize that experienced designers will be more adept at synthesis i.e., recombining information from prior highly-rated design submissions made by other designers, and by doing so will increase their probability of winning a contest.

To test my hypotheses, this study uses large-scale data from an online crowdsourcing platform for design tasks. I focus on design contests for logos, wherein contest holders create contests and invite designers to submit their solutions and the

winning solution is awarded a monetary reward by a contest holder. In particular, using deep learning and state-of-the-art image analysis techniques I examine how experienced designers incorporate information from highly-rated prior submissions from other designers within a contest.

I find that, surprisingly, less-experienced designers are, on average, similar to experienced designers in a number of ways. Both experienced designers as well as less-experienced designers are equally likely to wait for highly-rated submissions before making their own submission. Both types of designers are equally likely to incorporate information from highly-rated prior submissions within a contest. However, despite these similarities, I find that experienced designers are more likely to win a contest as compared to less-experienced designers.

In examining further, I find that there are crucial differences in how they incorporate information from these highly-rated prior submissions. I find that less-experienced designers' submissions are significantly more similar to the individual highly-rated prior submissions in a contest, which suggests a higher degree of imitation of individual highly-rated submissions by less-experienced designers. Interestingly, while experienced designers also incorporate information from prior submissions, their submissions are less similar to each of the individual prior highly-rated submissions as compared to the less-experienced designers.

Using state-of-the-art “deepfake” synthesis techniques, I further examine if experienced designers are more adept at such synthesis i.e., recombining information from different submissions, as suggested by prior research. I match designers by their position (i.e., order of first submission) in a contest and find that experienced designers

are indeed better at integrating information from multiple highly-rated prior submissions from other designers. I confirm this by using a neural style transfer analysis technique. With respect to the contest winning probability, overall, I find that while both leveraging information spillovers from individual highly-rated submissions and synthesizing information from multiple highly-rated submissions have an inverted U-shaped effect on the probability of winning a contest, the information synthesis has a significantly larger effect on the winning probability as compared to the effect of imitation of individual highly-rated submissions.

This study has a number of interesting theoretical and practical implications. While prior studies have attempted to speculate on what differentiates experienced designers from less-experienced designers in a variety of contexts, they have been limited by lack of granular data and techniques to help uncover the underlying mechanics. This study is among the first to explore the differences among experienced and less-experienced designers in leveraging information spillovers and effectively recombining information from prior submissions in crowdsourcing contests. Prior studies in other contexts and conducted using laboratory settings or field surveys, find that experienced users and, specifically, experienced designers are better at synthesizing or recombining information from several sources as compared to less-experienced users and designers (Björklund 2013, LaToza et al. 2015, Riedl and Seidel 2018). I confirm these prior findings in the new context of a crowdsourcing platform for design tasks using a large-scale unstructured design image dataset and by employing state-of-the-art deep learning methods. In doing so, this study contributes to the research streams that focus on the role of expertise in creative tasks. This study is also among the first to directly

measure information externalities in open contests, and, thus, it contributes to research focusing on information spillovers in crowd-based markets. Importantly, the findings of this study also contribute to the emerging research on recombinant innovations (Fleming 2001, Van den Bergh 2008, Frenken et al. 2012, Castaldi et al. 2015) by extending prior research into a new context of a crowdsourcing platform for design tasks. Studies on recombinant innovations in the fields of innovation and strategy claim that most innovations happen by recombining prior knowledge rather than by radically developing breakthrough innovations from scratch. I show that by recombining information spillovers from several prior highly-rated submissions made by other designers, experienced designers increase their probability of winning. This is consistent with the fact that prior highly-rated submissions incorporate a contest holder's feedback, and that ignoring this feedback by developing new design solutions from scratch would be risky for designers. Importantly, contest holders reward this synthesis or recombination of information spillovers by choosing these "synthesized" solutions as the winning solutions.

This study also has practical implications for the design of crowdsourcing platforms. Practitioners conduct experiments with both open and "closed" (or "blind") contests, as each of them has its own pros and cons. For instance, open contests make prior submissions available for other designers to view and that information visibility might help designers that enter later into a contest. The downside to such information disclosure is that some designers excessively copy prior solutions and that public feedback from a contest holder might reduce the variance of solutions (Wooten and Ulrich 2015). "Closed" or "blind" contests, on the other hand, do not disclose prior

submissions until the end of a contest. Contestants, consequently, cannot benefit from any associated feedback provided by the contest holder to select submissions. Thus, it is important to understand how the “open” nature of open contests might benefit or harm contest holders as well as designers by making information from focal contest submissions available for other designers to view.

The findings of this study show that information externalities from open contests can indeed benefit some designers as well as the contest holders. The findings show that while experienced designers leverage the information from highly-rated prior submissions from other designers in the contest, they do not necessarily imitate these submissions blindly. Rather, it is the experienced designers who are adept at synthesizing or recombining information from multiple highly-rated prior submissions, that are more likely to win a contest. Further, the less-experienced designers who are likely to imitate individual prior submissions more closely, are less likely to win a contest. Thus, the findings indicate that some of the concerns about the likelihood of imitation in open contests might be overblown. The findings also provide useful guidelines for designers in such marketplaces. Novice designers might be able to learn from observing the behaviors of experienced designers to obtain a more nuanced understanding of how their winning submissions are synthesized from prior highly-rated submissions. As for the contest holders, the findings highlight the importance of providing feedback that may play a crucial role in conveying their preferences for certain aspects or attributes of design submissions. The finding that winning submissions are more likely to incorporate elements of highly-rated prior submissions stands testimony to the importance of the visibility of the feedback provided by the

contest holders to early design submissions and their resulting benefits to the contest holders.

3.2 Related Research and Theoretical Underpinnings

I review prior related analytical and empirical research that mostly covers the optimal design of online crowdsourcing contests, designers' strategies and information spillovers and discuss the theoretical underpinnings of this study.

3.2.1 Design of Contests

Analytical research examining the design of crowdsourcing contests studies how the rules of contests affect designers' behaviors and contest outcomes. For instance, this stream of research explores the impacts of the number of prizes and the prize amounts (Archak and Sundararajan 2009, Gao et al. 2012, Ghosh and Kleinberg 2016), the optimal contest design that maximizes designers' effort (Körpeoğlu and Cho 2017) or the one that produces the highest revenue for a contest holder as well as for the whole platform (Wen and Lin 2016). Empirical research on contest design examines the optimal size of the reward and its relationship with participation and quality (Shao et al. 2012), the role and timing of feedback provided by contest holders (Wooten and Ulrich 2017, Jiang et al. 2018), the role of task design (Chen et al. 2014) and type of contests (i.e., open versus "blind", see Wooten and Ulrich 2015). For instance, Shao et al. (2012) find that higher rewards and longer duration attract more contest participants. Recent empirical research also explores the effect of task design and task description on the type of designers that a contest attracts and on the competition in the contest. For instance, Chen et al. (2014) study the role of project type (ideation or expertise)

and project complexity and find that ideation-based contests are very sensitive to monetary incentives, as compared to expertise-based contests. Jiang et al. (2019) examine the role of “conceptual objectives” and “execution guidelines” in problem specification and find that more conceptual objectives attract fewer contest participants, while more execution guidelines increase the trial effort by each designer.

Most research studies focus on analyzing either open (sequential) or “blind” (simultaneous) contests. Some recent empirical research explores the difference between open and “blind” contests. Research in this area examines designers’ selection of contests and designers’ efforts and their relationship to the type of contest. For instance, Jian et al. (2017) find that the maximum effort exerted by designers is higher in “blind” contests, while in open contests designers may alter their effort based on prior entries by other designers.

Continuing this line of research, this study examines how information spillovers from rival designers’ submissions are used by designers and how they impact designers’ outcomes. This study also leverages data available from both open and “blind” contests to study how experienced designers differ from less-experienced designers in leveraging information from highly-rated prior submissions.

3.2.2 Designers’ Strategies

Analytical research related to designers’ strategies focuses on the selection of contests by designers (DiPalantino and Vojnovic 2009, Segev 2020), on the response of designers to observed competition (Gross 2016) and on the effort that designers choose to exert (Körpeoğlu et al. 2017). Empirical research examines a variety of strategies adopted by designers, including which contests they choose to participate in, their

timing of entry, and how these are influenced by competition and the presence of other designers. DiPalantino et al. (2011) find that designers choose which contests to participate in depending on the prize ranges in accordance with their skill level. Shao et al. (2012) study crowdsourcing contest designers' reaction to competition intensity and find that lower competition intensity attracts more designers but not necessarily the more skilled designers. Zhang et al. (2019) find that competing in a contest with a “superstar” coder helps designers learn faster, increasing their probability of winning other contests.

Prior research also examines the timing of entry by designers in a contest. For instance, Yang et al. (2010) find that winners are more likely to submit early or later during the submission period but submitting in the middle of the contest lowers the likelihood of winning. They also find that strategic waiting to submit solutions increases the probability of winning.

This study continues this line of research with the goal of understanding designers' strategies in open crowdsourcing contests by focusing on information spillovers from prior submissions in open contests.

3.2.3 Information Spillovers

Information spillovers or information externalities are studied in many fields, such as economics (Choi et al. 2019), information systems (Hitt and Tambe 2006, Janze 2016, Krijestorac et al. 2017), marketing (Jing 2018) and finance (Chen et al. 2012, Wu et al. 2020). The main idea of information spillovers is that, by looking at other similar agents' behaviors, a focal agent can partially compensate for the lack of information about some important variables. In open crowdsourcing contests, while a contest

holder's preferences are revealed at the initial task specification stage, there is still a lot of uncertainty about what kind of submissions would meet a contest holder's requirements. However, the feedback (i.e., star rating) provided by a contest holder to early design submissions can be very useful to late movers in revealing valuable information about a contest holder's preferences. Later submissions that are successful in leveraging such information spillovers will likely have higher likelihood of being selected as winners by the contest holder.

Much of the research relating information spillovers is analytical due to the difficulty of measuring spillovers (Jing 2018, Choi et al. 2019). Prior empirical research studies spillovers from IT investments in offline settings (Hitt and Tambe 2006, Cheng and Nault 2012, Menon 2018). More recent research examines information spillovers in online settings. For example, Kwark et al. (2016) study the spillover effect of user-generated online product reviews of related substitutive and complementary products on the purchase of a focal product. The authors find that there is a negative spillover effect of online product reviews for substitutive products and a positive spillover effect for complementary products in the purchase of a focal product.

Some prior research examines information spillovers in crowdfunding markets. For example, researchers study the effect of funding spillovers from "blockbuster" projects to other projects in the same cluster and outside cluster and find evidence of positive spillovers (Kim et al. 2016). Kim and Viswanathan (2018) study information externalities in an online crowdfunding market and find that early investors' experience serves as an informational signal of quality for later investors. Senney (2019) finds that listings with early bidding in an online peer-to-peer loan market attract more lenders.

A few studies explore information spillovers in online crowdsourcing platforms. For instance, Deck and Kimbrough (2017) use a lab experiment to study how the prize allocation (“winner-take-all” or shared) and information disclosure about designers’ decisions and outcomes in prior periods (private or public) affect outcomes. The authors find that public disclosure encourages copycat behavior when the prize is shared, indicating that a “winner-take-all” contest is preferable to a shared prize contest in the case of public disclosure of performance. On the other hand, a “winner-take-all” approach is less preferable when outcomes from early attempts at innovation remain private, leading them to conclude that the optimal contest is a private information prize sharing contest. Hofstetter et al. (2020) study information spillovers in open crowdsourcing design contests and show that later-entering designers tend to copy prior highly-rated submissions from competing designers, but they also find that such copying behavior lowers the designers’ probability of winning. While this study provides insights into excessive imitation and associated penalties, it is still unclear whether and how prior submissions might have a positive spillover effect on later entrants, given that prior submissions, especially highly-rated submissions, reveal a contest holder’s preferences.

While some prior research examines information spillovers in crowdsourcing contests, there are still gaps in our understanding of how those spillovers are leveraged by contest participants and how experienced designers might process those spillovers differently compared to less-experienced designers. This study seeks to fill this gap by examining information spillovers at the individual level in an online crowdsourcing

marketplace and by exploring how such information spillovers are used by designers in open contests.

3.2.4 Experienced Designers and Recombinant Innovations

Prior research examines how experienced designers differ from the other designers. Researchers (Yang et al. 2010, Khasraghi and Aghaie 2014) find that past performance of designers is a good indicator of future winning probability. Archak (2010) finds that highly-rated designers take part in contests to deter other designers, while Boudreau et al. (2012) find that top-skilled competitors react to the presence of superstars with higher effort and performance. Ericsson et al. (2018) find that expert users are different in their problem-solving performance from novice users.

Research exploring problem-solving by designers provides more evidence that experienced designers are different from less-experienced designers. Of particular relevance to this study is research examining how experienced designers have the ability to leverage information from multiple sources and recombine them to create better solutions. For instance, Riedl and Seidel (2018) study innovation communities and find that more-experienced designers can successfully integrate signals about what is valued by the firm hosting the innovation community and by the community itself. These designers improve their performance by observing good examples from other designers and by synthesizing learnable signals from different sources. While experienced designers are more effective in representing design problems (Cross 2004), researchers (Atman et al. 1999, Popovic 2004) find that the more-experienced designers also gather and use more information related to the focal problem and are more active at identifying more information needs as well as more information sources.

Experienced designers tend to refer to past designs, and approach problem information more critically by questioning data and understanding data limitations and can differentiate between issues based on their importance (Ahmed et al. 2003). More importantly, Björklund (2013) finds that the experienced designers combine elements from different informational cues and integrate them more effectively and are able to determine which pieces of information are interconnected.

The concept of information integration or information synthesis is similar to the concept of recombinant innovations (Fleming 2001, Van den Bergh 2008, Frenken et al. 2012, Brynjolfsson and McAfee 2014, Castaldi et al. 2015, LaToza et al. 2015, Youn et al. 2015). According to the perspective on recombinant innovations, innovations do not involve coming up with something big and new, but instead recombining things that already exist (Brynjolfsson and McAfee 2014). While innovative breakthroughs are possible, they are much more resource-demanding and do not happen as often as compared to innovations that are created by combining existent and related knowledge (Castaldi et al. 2015). Frenken et al. (2012) compare innovations along a certain path (i.e., so-called “branching innovations”) with recombinant innovations, and argue that recombinant innovations speed up technological progress. Zhang et al. (2019) add to this discussion by introducing the concept of recombinant distance. Zhang et al. (2019) seek to understand how a focal firm’s distance from partners in a knowledge network affects recombinant innovation performance of a focal firm. The authors find an inverted U-shaped relationship between the focal firm's recombinant distance and its recombinant innovation performance and conclude that sources for recombining

knowledge should be sufficiently diverse to increase innovation performance, but not too diverse to inhibit such performance.

In the open innovation context, the “open” nature of contests and the feedback from the contest holder create an opportunity for designers to experiment with recombining different design elements (shapes, colors, typefaces etc.) from prior highly-rated design submissions within the contest, especially given that highly-rated prior submissions already incorporate design elements that are highly valued by the contest holder.

This study continues this line of research by examining how the strategies of experienced designers differ from those of novice designers in an online crowdsourcing marketplace. Specifically, this study uses state-of-the-art image analysis techniques and seeks to understand whether experienced designers resort to recombining design elements from prior submissions in crowdsourcing contests, and whether recombining design elements from prior designs benefits the designers and increases their probability of winning a contest.

3.2.5 Hypotheses

Based on the review of related research literature, I formulate hypotheses about possible directions of the effects.

Prior research mostly conducted in laboratory settings and by using surveys indicates that experienced users in general, and experienced designers specifically, are different from novice users and novice designers in their problem-solving skills (Ericsson et al. 2018) and in their information integration (or synthesis) capabilities (Björklund 2013, Riedl and Seidel 2018). Further, the perspective on recombinant innovations suggests that local search or “exploitation” (March 1991) is a process when

an inventor recombines from a familiar set of technology components. In contrast, distant search or “exploration” is the opposite case, when inventors experiment with completely new components or combinations. One can expect that local recombination would provide more certainty (because inventors would learn from past failures) and would be more successful, on average as compared to distant search (Fleming 2001). In the context of information spillovers in open crowdsourcing contests, this means that later-entering designers can potentially leverage early designs¹⁰ and use local “exploitation” strategies by recombining prior design elements into new design solutions. Additionally, inventors and designers, as they gain more experience, learn which combinations are more useful and less useful, and can use that combinatorial knowledge of optimal relationships among components to increase their inventive successes (Fleming 2001, LaToza et al. 2015). These arguments lead me to hypothesize that:

Hypothesis 1. Experienced designers will be more adept at recombining information from prior highly-rated submissions (made by other early-entering designers) as compared to less-experienced designers.

This strategy of recombining information from different informational cues is different from excessive imitation of individual submissions. Researchers (Fleming 2001, LaToza et al. 2015) find that recombining information can help create novel design solutions that incorporate several important design elements from different designs. In contrast to “branching innovations” that are related to technological

¹⁰ Later-entering designers can specifically leverage highly-rated early designs, since lower-rated designs would represent “failures” and designs that are not rated would represent uncertain designs with no feedback.

improvements along a certain path (in our case it is an individual improvement of a single prior highly-rated submission), recombinant innovations represent a fusion of several paths (Frenken et al. 2012), or, in our case, a synthesis of several highly-rated submissions made by other designers. In the first case of “branching innovations”, there is a risk of excessive imitation and associated penalties (Hofstetter et al. 2020), while in the case of recombinant innovations this risk is likely less severe. For instance, a designer might copy some part of a first highly-rated submission and another part of a second highly-rated submission to produce a novel design solution that might look differently and “farther away” from these individual highly-rated submissions (by maintaining the “recombinant distance”, as described in Zhang et al. 2019) but at the same time this design solution could incorporate valuable design elements from both submissions. In other words, using information spillovers from several highly-rated submissions may reduce the risk of penalty for excessive imitation, but at the same time allows incorporating valuable information from a contest holder’s feedback and preferences revealed through high star ratings assigned to prior submissions. These arguments lead me to hypothesize that:

Hypothesis 2. The recombination of information from multiple prior highly-rated submissions will provide experienced designers an edge in these open contests and increase their probability of winning a contest.

3.3 Research Context

3.3.1 Contest Description

This study uses a dataset from a large online platform for crowdsourcing design for logos, websites, apps etc. For the purpose of this study, I focus on the logo designs, as they represent the majority (about 75%) of designs on the focal platform. Clients (contest holders) submit a contest description and invite participants (designers) to submit their designs for a monetary award. During contests, contest holders provide numeric star ratings to logo submissions that they like or dislike. At the end of the “winner-take-all” contest, one designer is chosen by a contest holder as a winner.

To initiate a contest, a contest holder needs to do the following: The platform offers 120 design inspirations, from which a contest holder needs to choose at least 3 and at most 10 designs that a contest holder likes. This step is needed to evaluate style preferences of a contest holder. Image inspirations are offered from prior logos designed on the platform. This step could be skipped, but in my dataset about 89% of contest holders provide this information. A contest holder also provides a contest description including a contest title, name to use in a logo, slogan to incorporate in a logo, organization description and target audience, industry, other notes, and may provide image references. Image references are any images from external sources. Finally, a contest holder may choose a reward, the type of contest (open or “blind”), and contest duration (standard is 7 days, but this could be changed for an additional fee), and whether a prize is guaranteed or not. In a “blind” contest, prior submissions by other designers are not visible till the end of a contest.

Contests are normally run for 7 days, after which a contest holder should select a winner. Designers on the platform can view several contests and decide which one to participate in. Each contest has a description, and most contests (89%) provide image inspirations, while some contests (31%) also provide image references. Contest participants can view prior entries in open contests and numeric star feedback for those entries as well as numeric orders of entry. Upon clicking on a designer profile for those designers who made prior entries, designers can see other designers' experience (membership start date, number of contests won). Once a contest duration ends, a contest holder announces the winner.

3.4 Data and Methodology

I collect data for a period of 3 months in the Summer 2019. The sample includes 9,987 contests out of which 8,574 contests (85.85%) are open and the rest are “blind”. The total number of submissions is 626,979. I also collect a logo image for each submission as well as logo images for references and inspirations provided by contest holders.

Table 3.1 lists variables and descriptive statistics for each variable.

3.4.1 Structural Similarity

To calculate similarity variables *sim_ref*, *sim_insp*, *sim_high_star*, I use a structural similarity index (SSIM) algorithm (implemented in python in the scikit-image library¹¹) that is widely used in computer science research (Zhou Wang et al. 2004). This algorithm can be viewed as a quality measure of one of the images being

¹¹ https://scikit-image.org/docs/dev/auto_examples/transform/plot_ssim.html

compared. The second image is assumed to have a perfect quality¹². The structural similarity index has been shown to be superior to Mean Squared Error in terms of prediction of human perception of image fidelity and quality (Wang and Bovik 2009). The SSIM algorithm captures higher similarities of two identical images, even when one of the images is significantly altered. This is especially helpful for our case of logo comparisons in open contests where designers may observe prior logo submissions and leverage information from highly-rated logos submitted by other designers.

The SSIM index evaluates the similarities of three elements of the image patches: the similarity of the local patch luminances (brightness values), the similarity of the local patch contrasts, and the similarity of the local patch structures (Wang and Bovik 2009). The index is symmetric, i.e., gives the same metric if an image A is compared to an image B and vice versa. And it is bounded between -1 and 1, although in this study I find that all values are between 0 and 1, where 1 is perfect similarity. Examples of similar and dissimilar logos are shown in Figure 3.1 and Figure 3.2.

3.4.2 Empirical Models

The main dependent variable is a winner dummy indicating whether an image submission won a contest. To test how initial image references and image inspirations provided by a contest holder are used by designers and how these variables affect the probability of winning, I use independent variables similarity with references (`sim_ref`) and similarity with inspirations (`sim_insp`) described above and also use their squared terms to capture any non-linear effects. To test how information spillovers are used in

¹² <https://ece.uwaterloo.ca/~z70wang/research/ssim/>

contests, I use independent variable similarity of a focal submission with the most recent highly-rated submission by another designer (sim_high_star)¹³. To understand how experienced designers differ from novice designers in whether and how they leverage information provided by the contest holder as well as information from prior submissions, I use similarity with references (sim_ref), similarity with inspirations (sim_insp) and similarity of a focal submission with the most recent highly-rated submission by another designer (sim_high_star) as independent variables, and “experience” variable (as well as the “contests_won” variable as a robustness check, see Table 3.1) as the dependent variable in the model with contest fixed effects (see Table 3.4).

The general model for examining these effects is a logit model with contest fixed effects in the following form:

$$\text{Logit}(\text{Winning}_{ic}) = F(\text{sim_ref}_{ic}, \text{sim_ref_squared}_{ic}, \text{sim_insp}_{ic}, \text{sim_insp_squared}_{ic}, \text{sim_high_star}_{ic}, \text{sim_high_star_squared}_{ic}, \text{sub_order_log}_{ic}, \text{experience_log}_{ic}, \text{star}_{ic}) + \alpha_c + \varepsilon_{ic} \quad (3.1)$$

In the above equation (3.1) i represents submission in a contest c , α_c represents a contest fixed effect. The unit of analysis is thus a submission i in a contest c .

To capture the timing of entry behavior by each designer, I also measure $\text{first_sub_order}_{dc}$, the order of the first submission by each designer d in a contest c . This variable is used in a separate model (where unit of analysis is a designer in a

¹³ I also examine the variables representing similarity with the second most recent highly-rated submission, similarity with the third most recent highly-rated submission etc. I find that each focal submission, on average, is closer to the most recent prior highly-rated submission, so I use this variable in the main analyses. Nevertheless, using second, third, fourth (and so on) most recent highly-rated submissions provides consistent results as compared with using the similarity with the “first” most-recent highly-rated submission.

contest) to test whether designers who enter earlier or later are more likely to win a contest (see the section “Effects of Timing of Entry and Designer Experience” below).

The contest fixed effects model (3.1) evaluates differences among designers in a contest. Since some designers make multiple submissions per contest, as a robustness check I also use a more restrictive model with designer and contest fixed effects. This model requires that a designer has made more than one submission per contest. But this model is expected to control for heterogeneity among different designers. I report the results for the contest fixed effects model (3.1) and find the results to be consistent with the contest and designer fixed effects model (see Appendix, Table A4, Model 4).

3.5 Results

3.5.1 Effect of Information Provided by the Contest Holder

I first evaluate the effects of initial information provided by the contest holder, specifically image references and image inspirations on the probability of winning a contest. Similarity with image references is correlated with similarity with image inspirations (correlation is 0.51). In addition, I use squared terms of those variables, which creates additional correlations between simple and squared terms. Thus, I evaluate the effects of similarity with references (sim_ref) and similarity with inspirations (sim_insp) separately (model 1 in Table 3.2), while controlling for the experience of a user making a submission, the star rating of a submission and the order of a submission. Model 1 shows that similarity with references as well as similarity with inspirations have an inverted U-shaped effect on the probability of winning a contest, with the inspirations having higher effect size than references.

As a robustness check, I include all similarities in one model (Model 3 in Table 3.2) to see which similarities are still significant, controlling for other similarities. In the Model 3 similarity with references becomes non-significant, while similarity with inspirations is significant and has an inverted U-shaped effect on the probability of winning a contest. Additionally, as a robustness check, I restrict the sample to contests with only one type of images provided by a contest holder, either image inspirations or image references. I find the results to be consistent with the Model 3 (see Appendix, Table A4, Model 1 and Model 2).

As an additional robustness check I include designer and contest fixed effects. This model (see Appendix, Table A4, Model 4) provides very similar estimates as the model with the contest fixed effects.

3.5.2 Effect of Information Spillovers from Prior Submissions

Next, I evaluate the effect of similarity of a focal submission with the most recent prior highly-rated submission made by other designers, on the focal submission's probability of winning a contest. Again, similarity with prior most recent highly-rated submissions is correlated with similarity with image inspirations (correlation 0.577) and correlated with similarity with image references (correlation 0.396). Additionally, I use squared terms of this variable, which creates additional correlations between simple and squared terms. Thus, I evaluate the effects of similarity with the prior most recent highly-rated submissions (`sim_high_star`) separately from other two similarities (model 1 in Table 3.2), while also controlling for the experience of a user making a submission, the star rating of a submission and the order of a submission.

Model 1 in Table 3.2 shows that similarity with prior most recent highly-rated submissions has an inverted U-shaped effect on the probability of winning a contest. As a robustness check, I combine all similarities in one model (Model 3 in Table 3.2) to see which similarities are still significant controlling for other similarities. In the Model 3 similarity with references becomes non-significant, while similarity with inspirations and similarity with prior most recent highly-rated submissions are significant and have an inverted U-shaped effect on the probability of winning a contest. Additionally, I restrict the sample to contests where image inspirations/references are not provided by contest holders and find the results for prior highly-rated submissions to be consistent with the models 1 and 3 (see Appendix, Table A4, Model 3).

As with the prior analysis, an additional robustness check involves using designer and contest fixed effects. This model (see Appendix, Table A4, Model 4) provides very similar estimates as the model with contest fixed effects.

3.5.3 Effects of Timing of Entry and Designer Experience

In this section I evaluate the impact of timing of entry on the probability to win a contest as well as whether experienced designers enter contests earlier or later than less-experienced designers.

Table 3.2 shows the results for `sub_order_log` for open contests in Models 1 and 3. The effect of submission order is positive and significant, which means that later entries have higher probability to win a contest. This result is consistent if I use the order of first submission (`first_sub_order`) by each designer.

Table 3.3 shows that experienced designers are more likely to make their first submission later than less-experienced designers.

I also check whether experienced designers are more likely to enter open contests after a prior highly-rated submission has been made by another designer. Surprisingly, I find that, even though experienced designers are more likely to enter a contest later compared to less-experienced designers, experienced designers do not differ from less-experienced designers in their likelihood of entering a contest after a prior highly-rated submission has been made by another designer.

3.5.4 How do Experienced Designers Differ from Other Designers?

Table 3.2 (Model 1 and Model 3) shows that experienced designers are more likely to win a contest compared to less-experienced designers. However, given that there are multiple images, those provided by the contest holder, as well as those from highly-rated prior submissions, I first rank the images that are closest in similarity to the designers' submissions. I find that, on average, the designers' submissions are closest to the two highly-rated prior submissions from other designers. In Table 3.1 the mean similarity of a focal submission with the most recent prior highly-rated submission by another designer is higher than the mean similarity with inspirations and the mean similarity with references. Additionally, I check similarities of a focal submission with multiple prior highly-rated submissions and find that the mean similarity of each submission with the most recent two prior highly-rated submissions by other designers is higher than the similarity with any other image provided by the contest holder or other designers. Similarity of a focal submission with the more distant (third prior and earlier) submissions decreases with the distance, which means that the submissions that

are farther away from a focal submission are less likely to be copied by designers. Thus, I confirm that the two closest images for each submission are the two most recent prior highly-rated submissions. In further examining the two subgroups - experienced designers and less-experienced designers – I find that both experienced designers’ submissions as well as less-experienced designers’ submissions are closest to the two highly-rated prior submissions in comparison to the image inspirations and references provided by the contest holder. This is also likely to be the case as the highly-rated prior submissions already incorporate information from the image inspirations and references provided by the contest holder.

I then examine the differences in the similarities between experienced and less-experienced designers’ submissions and the highly-rated prior submissions. As seen in Table 3.4, I find that experienced designers are less likely to excessively copy inspirations and individual prior most recent highly-rated submissions as compared to less-experienced designers. This is also the case for image inspirations in the case of “blind” contests where designers do not have access to prior submissions.

Prior research indicates that experienced designers not only search for more information in design tasks, but are also able to better integrate that information in their designs (Björklund 2013, LaToza et al. 2015). Given that, for both experienced as well as less-experienced designers, the submissions are closest in similarity to highly-rated prior submissions, I focus on how these designers leverage information spillovers from these prior submissions from other designers within a contest.

To examine how experienced designers differ from less-experienced designers in recombining or synthesizing information from prior submissions, I match experienced

and less-experienced designers by their position (i.e., order of first submission) in a contest. The matching is done in the following manner: first, I ensure that the matched users enter a contest after two highly-rated submissions have been made in a contest by prior designers; second, I ensure that the first entry of matched designers does not have a star rating, i.e., the submission should have been made after two highly-rated submissions but before other submissions with any star rating. Next, I ensure that there is more than one submission without a star rating made by designers (to be matched) after the first two highly-rated submissions. This positioning is very restrictive, but I want to capture the difference of experienced and less-experienced designers in how they leverage such information.

After the matching, I then examine whether and how designers recombine information from the two highly-rated prior submissions¹⁴ within a contest. I use a state-of-the-art deep learning algorithm for image synthesis (Karras et al. 2019) implemented in python¹⁵. The algorithm was first developed for producing “deepfake” images based on the training sample. It uses a generative adversarial network (styleGAN) to produce fake images and combine or synthesize images by mixing the style of images being synthesized. I have trained the algorithm for several weeks on a random set of 23,000 logo images and have been able to get the fid50k metric close to the value of 25¹⁶. This metric compares real logo images and fake logo images and

¹⁴ While there is a possibility that designers might copy design elements from more than two prior highly-rated submissions made by other designers, algorithmic merging of three design images is technically much more sophisticated, and it will create synthesized images that look “blurred”, making hard to see whether a synthesized design image represents an integration of three images. Thus, I keep the merging analysis to combining two images for simplicity of validation.

¹⁵ <https://github.com/NVlabs/stylegan>

¹⁶ While the ideal value would be less than 10, that is the best value I could achieve.

measures how far real logos are from fake logos. But this method can only synthesize fake logos generated by the styleGAN network. To synthesize my own images from the matched set, according to the method proposed by Dmitriy Nikitko¹⁷, I embed the focal logos with high stars from the matched sample into the latent space of the prior styleGAN network and create a synthesis of two images with high stars for 1,000 contests. Thus, in total, I merge 1,000 pairs of high-star logos to produce 1,000 synthesized logos that use styles from both highly-rated images in the proportion of 50/50. The reason for choosing this proportion is related to the finding that, on average, a focal logo submission is close to two most recent highly-rated submissions in a similar proportion. I have also tried synthesizing images with the following proportions – 70/30, 30/70, 10/90 and 90/10. Interestingly, the synthesis in the proportions of 70/30 and 30/70 provide results consistent with the synthesis in the proportion 50/50. But the synthesis in the proportions of 90/10 and 10/90 provides results consistent with the Table 3.4, which means that this type of synthesis (where 90% is coming from one image) is close to imitation of an individual highly-rated submission. Specifically, experienced designers' submissions are less similar to the image that is a synthesis of two images in proportions of 90/10 or 10/90, as compared to less-experienced designers' submissions (see Table A5 in Appendix).

Examples of two original logo images (images A and B) and their synthesis (image C) are shown in Figure 3.3.

¹⁷ <https://github.com/Puzer/stylegan-encoder>

As a robustness check, I use a second method to synthesize images - a neural style transfer based on deep neural network¹⁸. This method is based on the paper by Gatys et al. (2015) and improvements suggested by Novak and Nikulin (2016). Example of a synthesized image produced by this algorithm is shown in Figure 3.3, where the image D is a synthesis of images A and B using the neural style transfer method. This second method produces a similarity metric that is highly correlated with the first method (correlation is 0.664). Examples of source images A and B, synthesized images with the neural style transfer and corresponding winning images are shown on Figure 3.4.

Next, I compare the similarity of matched experienced and less-experienced designers' submissions with the synthesized images. I find that experienced designers' submissions are more similar to the synthesized images (merged in the proportions of 50/50 or 70/30 or 30/70) as compared to the matched less-experienced designers. For instance, if experience is increased by 1 percent, similarity with the synthesized images increases by 0.031%. Estimates from the second neural style transfer method are consistent and provide the value 0.0347, which means that if experience is increased by 1 percent, similarity with the synthesized images increases by 0.0347% (see Table A5 in Appendix). In other words, if experience increases by 100%, I can expect that a submission of such designer will be closer to the synthesis of two prior most recent highly-rated submissions by 3.47%. It is pertinent to note that the variable "contests_won" (for the number of prior contests won by a designer, see Table 3.1) provides similar and consistent estimates as the variable "experience" (which indicates the number of days since joining the platform by a designer, see Table 3.1).

¹⁸ <https://github.com/titu1994/Neural-Style-Transfer>

Specifically, if the number of contests won by a designer increases by 100%, I can expect that a submission of such designer will be closer to the synthesis of two prior most recent highly-rated submissions by 5.35% (see Table A5 in Appendix). Thus, the *hypothesis 1* is supported, which means that experienced designers are more adept at recombining information from highly-rated prior submissions made by other designers as compared to less-experienced designers.

Additionally, as shown in Table 3.2, similarity with synthesized images has an inverted U-shaped relationship with the probability of winning, and the effect size is much larger than other similarities. Thus, the *hypothesis 2* is supported, which means that this recombination of information from several prior highly-rated submissions provides experienced designers an edge in open contests and increases their probability of winning a contest as compared to less-experienced designers that tend to “blindly” imitate individual highly-rated submissions.

3.5.5 Comparison of Open and “Blind” Contests

Since “blind” contests are different from open contests in their information visibility, I check if information spillovers are not important in such contests. To do that, I compare the findings of open contests with “blind” contests. Both open and “blind” contests have initial project information from a contest holder available to designers (including project description, image references and image inspirations). However, in a “blind” contest, designers are only able to view their own entries. While they are able to see any feedback (the star ratings) provided by the contest holder for prior submissions, the submissions themselves are not visible to other designers. Thus, in “blind” contests, any potential information from prior highly-rated submissions is unlikely to affect a

designer's winning probability for both experienced as well as less-experienced designers.

Model 2 in Table 3.2 for “blind” contests shows that similarity with references and similarity with inspirations have an inverted U-shaped relationship with the probability of winning, while similarity with prior highly-rated most recent submissions does not have a significant effect on the probability of winning, as expected.

Experienced designers have a higher probability of winning a contest in both open and “blind” contests (Models 1 and 2 in Table 3.2). Experienced designers are more likely to enter both open and “blind” contests later as compared to less-experienced designers (Table 3.3). But experienced designers, as compared to less-experienced designers, are less likely to excessively copy individual prior highly-rated submissions and inspirations in open contests, while in “blind” contests experienced designers are less likely to excessively copy inspirations compared to less-experienced designers (Table 3.4).

3.5.6 Robustness Check for a Potential Confounding Effect of Task Textual Descriptions

An alternative explanation for the observed effects could be that some designers are better integrating information in task textual descriptions provided by a contest holder. If they do so, later submissions might be more similar to prior highly-rated submissions because those later submissions and prior highly-rated submissions both leverage more information from a prior task textual description provided by a contest holder. Thus, to control for that, I restrict the sample to 935 contests where task descriptions are minimal and do not provide any useful information that could help increase the probability of

winning a contest. I find the results to be consistent with the main findings of this study (see Appendix, Table A4, Model 5).

3.6 Implications and Conclusion

Emergence of crowdsourcing marketplaces, including markets for labor and innovation, as an alternative solution to traditional outsourcing, has opened new opportunities for businesses. Many of these marketplaces use open contests to encourage more participation and competition. But the open nature of contests might promote strategic behaviors among participants, with participants timing their entry into contests to benefit from information spillovers from prior participants' submissions in the contest. These strategic behaviors might be different for designers with different levels of experience as indicated by prior research.

This study seeks to explore strategic behaviors of designers in an online crowdsourcing marketplace for logo design with the focus on how experienced designers leverage initial information signals from a contest holder's image references/inspirations and information spillovers from prior highly-rated submissions as compared to less-experienced designers. This study also seeks to explore the effect of those references/inspirations and spillovers on contest outcomes.

The findings of this study have important theoretical and practical implications. First, prior research in contest design examines the impact of task textual descriptions and their relationships with designers' behaviors. This study extends this line of research and is among the first to evaluate informational signals from images provided by contest holders as part of the task assignment. I find that the relationship between the similarities of images provided by the contest holder as well as by prior designers

and the focal designers' submissions have an inverted U-shaped relationship with the probability of winning. In other words, while some degree of similarity is beneficial, being too similar to these other images decreases the likelihood of winning. Interestingly, I find that less-experienced designers' submissions are more similar to these individual images, particularly those from highly-rated prior submissions, as compared to the submissions by experienced designers.

While prior research examines information spillovers from IT investments in offline settings and from user-generated content in online settings, more recent research explores information spillovers in crowd-based markets. This study is among the first to directly measure information spillovers using state-of-the-art image analysis techniques in a crowdsourcing platform for design tasks.

My review of prior research literature also indicates that there is a lack of empirical research related to the role of expertise in the usage of information signals and information spillovers in crowdsourcing contests. Prior research uses surveys of designers and laboratory experiments and finds that experienced designers (in other contexts) leverage more information from a design task assignment and can better integrate interconnected information from multiple sources (Björklund 2013). The findings of this study provide more granular insights into the differences between experienced and less-experienced designers. The finding that experienced designers are more adept at recombining information from multiple prior highly-rated submissions, sheds light on what differentiates experienced designers from less-experienced designers despite having access to the same set of information. Additionally, this finding contributes to research related to recombinant innovations by extending prior

research into a new context of a crowdsourcing platform for design tasks. I show that experienced designers can better recombine, or synthesize, several prior highly-rated submissions as compared to less-experienced designers, which is consistent with the concept of recombinant innovations. This approach is superior to “branching innovations”, when designers just improve individual prior highly-rated submissions, because in the latter case there is a risk of being penalized by a contest holder for imitation of prior submissions. In contrast, in the case of recombinant innovations, if designers keep some “recombinant distance” from other submissions (Zhang et al. 2019), the risk of being penalized for imitation is lower.

This study also has valuable managerial implications. First, the findings of this study could be used to better understand the role of information signals and spillovers from other images and the relationship of those informational elements with contest outcomes in open and “blind” contests. These insights can be useful for managers in improving the design of contests. While in “blind” contests information spillovers from prior submissions are not present due to visibility restrictions, in open contests those spillovers have interesting effects when excessive imitation of a single highly-rated submission is penalized by a contest holder, while the recombination of spillovers from multiple highly-rated submissions is beneficial for both contest holders and designers. Next, the findings related to designers’ expertise can provide guidance to less-experienced designers by helping understand how experienced designers integrate information from prior submissions to improve their likelihood of winning contests. Additionally, the findings highlight the important role of feedback for contest holders in open contests. Contest holders reveal their preferences for certain design elements

(shapes, colors, typefaces etc.) by providing star ratings to design submissions. The visibility of highly-rated prior designs gives valuable information to other designers which helps them to recombine or synthesize certain design elements and create solutions that are more beneficial for contest holders. Finally, managers might introduce better mechanisms protecting original submissions from excessive copying taking into consideration the finding that some copying is tolerated by contest holders and even helps increase the probability of winning a contest. Excessive copying can deter early entrants and platform managers could devise appropriate incentives for early entrants and compensate them for the positive externalities that they create for the later entrants, which eventually benefit the contest holders as well.

This study is not without limitations. First, while I show the effects of information spillovers using image similarity and image synthesis algorithms, I have not explored which specific components of an image tend to be copied or tend to be avoided by designers. Future research might decompose images into various elements to check if some parts of an image are especially helpful for winning a contest.

Next, in this study I do not have access to demographic information for designers. Future research might explore heterogeneous effects of leveraging information signals and spillovers depending on a specific demographic variable of interest such as age and gender.

Finally, I explore only one type of design – logo design, which represents majority of the tasks on the crowdsourcing platform. Future research might investigate other design assignments, such as design of a website or design of an app to see if findings depend on the type of a product being designed.

Chapter 4: Threatened by AI: Analyzing Users' Responses to the Introduction of AI in a Crowd-sourcing Platform

4.1 Introduction

With the rapid deployment of new machine learning and artificial intelligence (AI) solutions, these systems increasingly compete with human employees (Frey and Osborne 2017). Over the centuries a plethora of technologies have enabled the automation of routine tasks. These technologies, from steam engines to industrial machinery, have proven superior to humans on a variety of dimensions such as *power*, *speed*, *productivity*, *quality*, *accuracy*, *reliability*, *durability*, and often, in *cost* (Autor 2015, Chui et al. 2016). An analysis of 2,000 work activities across 800 occupations by McKinsey Global Institute finds that 60% of all occupations have at least 30% of activities that could be automated (Manyika et al. 2017). For businesses, automation has a lot of advantages that directly affect the bottom line which had led to a rapid adoption of these technologies by many organizations (Manyika et al. 2017). Given the superiority of these technologies in their specific tasks, they have quickly replaced humans that have traditionally performed those tasks and human workers have had to quit and switch to other “less automatable” tasks (Dixon et al. 2019).

Interestingly, AI solutions and technologies go beyond just automation of routine tasks as they are able to “learn” from past data to provide novel solutions. Increasingly, such AI systems are being deployed for creative tasks – which have largely been the domain of humans. In the case of creative tasks, faced with the threat of competition from AI, humans can either quit those tasks that the AI systems perform, or can compete with the AI systems by leveraging such qualities as imagination, creativity

and emotional expression (Hertzmänn 2018, Lu et al. 2020). There is very little systematic understanding of users' behaviors and strategies in response to the adoption of AI systems. In addition, studies examining the impacts of adoption of AI technologies have largely focused on their adoption at the level of industries, geographic regions, or within firms (Acemoglu and Restrepo 2017, Graetz and Michaels 2018, Mann and Püttmann 2018, Dixon et al. 2019). However, there are hardly any studies that examine the introduction of AI in decentralized marketplaces, or studies that examine how individuals respond to the adoption of AI systems for creative tasks.

This study seeks to fill this gap by studying a crowdsourcing platform for design tasks that has introduced an artificial intelligence system for logo design (AI logo maker) that can compete for design tasks with the human designers on the same platform. Access to large-scale granular data on designers' behaviors before and after the adoption of the AI system provides me an opportunity to study designers' heterogenous responses to an exogenous change - the introduction of an AI system for logo design by the platform. In this study I seek to answer the following research questions:

How do contest participants (designers) respond to the introduction of the AI system for logo design tasks?

How do the behaviors of successful contest participants differ from the behaviors of other participants upon the introduction of the AI system?

I study a crowdsourcing platform for design tasks that introduced an Artificial Intelligence system for logo design at the beginning of April 2018. I collect

comprehensive data on all design contests with human designers for the periods of 6 months before the AI launch and 6 months after the AI launch. The dataset includes 5,737 contests and the complete history of participation for 9,280 designers, which includes a total of 425,475 design submissions. I track the contestants who were available in both periods (before and after the AI launch) and compare the behaviors of successful and unsuccessful contestants to understand their responses to the introduction of the AI system. To understand whether and how contest participants change their design submissions in response to the AI system, I measure design image emotional content and complexity, since those variables have been shown (in prior research in psychology and marketing) to affect aesthetic perception of art and design images as described below. I use the SentiBank deep learning model (Borth et al. 2013) to build neural network models for measuring the emotional content of each design submission and use a spatial information complexity method to measure design complexity. I then examine whether and how the emotional content and complexity of design submissions affect the likelihood of winning a contest and how these differ for successful and unsuccessful contestants before and after the introduction of the AI system.

My identification strategy involves a “before-after” analysis for the same set of users active in contests on the platform both before and after the AI introduction. To support a causal identification of the effects, I also use propensity score matching techniques (PSM) as well as a difference-in-differences method (Smith and Todd 2005) to compare the effects in the treatment groups (i.e., users affected by the introduction

of the AI logo maker) with the effects in the control group (i.e., users participating in other non-logo contests both before and after the introduction of the AI logo maker).

In previewing the results, I find that the AI system cannibalizes lower-tier less-complex logo-design contests that have a lower award amount while it has no significant impact on the number of available contests in the higher-tier or in other (non-logo) design categories. This partial cannibalization shows that some clients preferred using the AI system instead of running contests with human designers. As for designers, I find that designers can be grouped into three categories based on their participation in contests, and that they respond differently to the introduction of the AI system. Designers who participated primarily in lower-tier logo-design contests either leave the platform or continue to participate in the lower-tier logo-design contests even after the introduction of the AI system. I term these focused designers. On the other hand, designers who had earlier participated in both lower-tier as well as higher-tier logo contests switch to higher-tier logo contests (I term these cross-tier designers), while designers who had prior exposure to contests in other categories for non-logo design are more likely to switch to participating in these non-logo design contests after the introduction of the AI system (I term these cross-category designers).

In examining how the behaviors of successful contestants differ from the behaviors of the other contestants after the introduction of the AI system, I find that in contrast to the unsuccessful contestants who increase the number of contests they participate in (by 13-15% depending on a group of designers), the successful contestants substantially increase the number of re-submissions (by 30-60% depending on a group of designers) within a contest as compared to the period before the introduction of the

AI system. Further, I also find that with an increase in the number of re-submissions by successful contestants in all three groups of designers – focused designers, cross-tier designers, and cross-category designers – there is a concomitant significant increase in the emotional content (the effect size is between 6.5% and 14.37% depending on a group of designers) as well as the complexity (the effect size is between 7.5% and 17.1% depending on a group of designers) of their design submissions after the introduction of the AI system. On the other hand, there is no significant change in the emotional content or complexity of the design submissions by the unsuccessful contestants when comparing the periods before and after the introduction of the AI system.

The findings show that the successful contestants behave in line with the findings of well-established research on the key factors that drive aesthetic experience in design. Seminal work in the psychology of aesthetic experience by Berlyne (1974) finds that two interrelated constructs affect human aesthetic experience – complexity and emotions. More recent research (Marin et al. 2016) finds that complexity and emotions are positively associated in creating an aesthetic experience, and that complexity is positively associated with beauty. More specifically, with respect to logo designs, Grinsven and Das (2016) show that logo design complexity positively affects long-term brand recognition and brand attitude in the case of repeated exposures, while De Marchis et al. (2018) find that emotions expressed in logos are strongly associated with the aesthetic attraction of logos. I find that the successful contestants focus on improving the emotional content and complexity of their design submissions by improving upon their original submission to a contest. Interestingly, this is in contrast

to most other contestants who choose to hedge their bets by making submissions to multiple contests. The focus on improving emotional content and complexity of design submissions by winners is also meaningful considering the current limitations of AI systems. Research on the limitations of AI shows that humans are better than AI systems when it comes to qualities such as creativity, imagination and emotions, in general (Braga and Logan 2017), and more so in the case of applications involving design and art, specifically (Hertzmänn 2018, Mazzone and Elgammal 2019).

The findings of this study have important theoretical implications and contribute to the nascent research stream related to employment effects of AI (Dixon et al. 2019) and employees' reactions to the competing AI systems (Lu et al. 2020). Prior research on AI effects on employees has been mostly conducted at the industry level and at the firm level (Acemoglu and Restrepo 2017, Dixon et al. 2019). This stream of research finds that overall effects of AI are negative at the industry level, but positive at the firm level for most workers except for managers who are more likely to be displaced by AI systems. This study is among the first to extend this line of research to a decentralized crowdsourcing platform where participants are free to choose how they respond to an exogenous shock on the platform. This study goes further to shed light on how successful designers are different from the others in how they respond to the introduction of the AI system on the platform and contributes to the emerging research that seeks to understand the effects of AI systems in business setting. Additionally, existing research (Wolbring and Yumakulov 2014, Li et al. 2019) that explores the responses of employees to the introduction of AI systems, has been mostly conducted using surveys of employees, and the findings have been context specific. This study is

among the first to use granular and longitudinal data to examine users' responses to the introduction of an AI system, and by doing so contributes to this nascent stream of empirical research in this area. The findings also highlight the differences in individuals' responses to the threat of competition from AI as compared to prior technologies which have essentially replaced humans performing those tasks.

Prior work in psychology (Berlyne 1974, Marin et al. 2016) and marketing (Pieters et al. 2010, Grinsven and Das 2016, De Marchis et al. 2018) has identified the role of emotions and complexity in driving the aesthetic perception of design in general as well as logo design in particular. This study is among the first to leverage large-scale granular data as well as state-of-the-art image analytics techniques, to empirically test these theories in an online crowdsourcing platform for design tasks.

The findings of this study also have important practical implications. Importantly, platform providers can use the findings to better evaluate the impacts of AI systems on contests and designers' behaviors. Understanding how different groups of users respond to the launch of an AI system can help market providers to design relevant pricing and marketing strategies to optimize performance and revenue from both sources – AI as well as human designers. A more nuanced understanding of the capabilities and limitations of AI systems relative to those of expert human designers can help platform providers recommend specific guidelines for contest holders as well as contest participants to improve outcomes. This could also pave the way for hybrid contests or hybrid solutions that leverage the capabilities of both the AI system as well as human experts.

This paper is organized as follows. Section 4.2 reviews related research. Research context and data are discussed in Section 4.3. Section 4.4 describes related methodology. Results are reported in Section 4.5. Finally, Section 4.6 discusses the main findings and implications.

4.2 Related Research

4.2.1 Impacts of AI on Organizations

Recent research studies the effects of adoption of AI systems on employment and skill composition (Acemoglu and Restrepo 2017, Graetz and Michaels 2018, Mann and Püttmann 2018, Dixon et al. 2019, Webb 2019). While some of this research examines the impacts of AI at the industry and geographic region levels, more recent research studies relevant outcomes at the firm level. At the industry level, the effects of AI on employment are mostly negative, while at the firm level they are more nuanced. For example, Dixon et al. (2019) find that, at the firm level, the employment effects of AI adoption are positive. The authors also find that, surprisingly, the adoption of AI is associated with the displacement of managers. Additionally, while prior research indicates that new technologies might replace lower-skilled workers (Graetz and Michaels 2018), Webb (2019) compares the text of patents with the text of job descriptions and finds that in contrast to software and robots, AI is directed at high-skilled tasks.

This study continues the above stream of research and seeks to understand how AI systems impact design contests and individual contestants on a crowdsourcing platform. Crowdsourcing platforms are different from other firms adopting AI systems,

since the participants on such platforms are freelancers. Thus, in contrast to prior work, this study seeks to investigate the impact of the adoption of the AI system on the users' behaviors and strategies on such platforms where participants are free to leave the platform at any time or switch to other available jobs on the platform.

Another related stream of research examines employees' responses to the adoption of AI systems, including their perceptions concerning job security and increased pressures to enhance skills and competences (Nam 2019, Lu et al. 2020). This stream of research finds that the responses of employees to AI adoption might be context specific. For instance, Wolbring and Yumakulov (2014) study the perceptions of smart AI robots by workers in disability care and find that workers do not feel threatened as they believe that these AI robots cannot replace human touch, human interaction, or emotional companionship. In contrast, Li et al. (2019) find that in the hospitality industry, hotel employees are more likely to quit if they are aware of the implementation of AI and robotic platforms in their organization. However, this decision is moderated by perceived organizational support. Similarly, Brougham and Haar (2018) find that higher awareness about AI applications is negatively related to organizational commitment and career satisfaction, and positively related to turnover intentions, cynicism, and depression. Most of these studies are conducted using surveys of employees.

This study extends this line of research and seeks to understand users' responses to the adoption of the AI system in a decentralized crowdsourcing platform for design tasks. What differentiates this research setting from prior research is that I can observe users' choices in response to the introduction of the AI system. More importantly, given

the decentralized nature of the platform and the diversity of participants, this study focuses on understanding how heterogeneous users respond to the introduction of an AI system that is a direct potential competitor for design tasks on the platform.

4.2.2 Contestants' Behaviors and Strategies on Crowdsourcing Platforms

This study builds on prior research that examines strategies and behaviors of contest participants in crowdsourcing platforms. Prior literature focuses on the behaviors and strategies of contest participants relating to responses to different project types and task specifications (Chen et al. 2014), to the choice of contests to participate in, and to the number of submissions within a contest (DiPalantino et al. 2011, Bockstedt et al. 2016). Chen et al. (2014) find that more-complex tasks typically attract fewer contestants, while Jiang et al. (2019) examine problem specifications and find that more conceptual objectives attract fewer contestants. In examining contestants' preferences for contests to participate in, DiPalantino et al. (2011) find that contestants choose contests depending on specific award ranges that correspond to their skill level, while Bockstedt et al. (2016) find that the number of submissions has a curvilinear relationship with the probability of a success in a contest. This study adds to this stream of research by examining how the introduction of the AI system affects the choice of contests by participants, and whether contestants change their behaviors after the AI launch.

A second related stream of research in this domain focuses on the differences in behaviors between successful contest participants and others, and finds that successful contest participants are typically more experienced (Khasraghi and Aghaie 2014), and they are strategic about timing their submissions (Yang et al. 2010). Yang et al. (2010) find that successful contest participants prefer submitting at the beginning or at the end

of a contest, while Archak (2010) finds that successful top contest participants might enter contests earlier to “deter” entries from other participants. Boudreau et al. (2012) also study successful participants in online crowdsourcing contests and find that successful top-skilled participants increase their effort in the presence of other top-skilled participants.

This study continues this stream of research and seeks to understand strategies and behaviors of successful contest participants as a response to the introduction of the AI system, as well as whether, and how, they differ from the responses of other contest participants.

4.2.3 AI Systems and Prior Technologies

Prior studies (Makridakis 2017, Brynjolfsson and McAfee 2014) have identified three distinct periods of technological evolution, beginning with the “industrial revolution” characterized by the domination of mechanical technologies ranging from steam engines to cars. These technologies that were superior in power and speed were primarily used for substituting *routine manual* tasks such as rowing, lifting objects or moving/walking etc. This period was followed by the “digital revolution” which started with the invention of the computer in 1946 and continues with the widespread usage of personal computers, smart phones, and networked devices. These digital technologies have proven superior to humans in *productivity, quality, accuracy, reliability, durability, and often in cost*, and have rapidly substituted humans in the performance of *standardized mental tasks* (Brynjolfsson and McAfee 2014, Autor 2015, Chui et al. 2016). The ongoing “AI revolution”, starting with neural net devices in the 1990s, and the development of more recent applications of computer vision and speech

recognition, is characterized by technologies that seek to mimic human brain power and cognitive abilities (Wang and Siau 2019) and could potentially perform all *mental* tasks (Makridakis 2017) including creative tasks such as design and art generation. Specifically, AI technologies are seen as distinctly different from prior generations of technologies in that they can learn and update using data such as numeric data as well as text, audio and video (Huang et al. 2019). These AI systems increasingly perform or simulate non-routine tasks requiring “tacit” knowledge by learning from prior successful examples of those tasks and using a process of exposure, training, and reinforcement (Autor 2015). While the AI systems cannot yet acquire the “tacit” knowledge that humans possess, they nevertheless can bypass this “tacit” knowledge requirement by being trained on millions of successful examples of tasks being simulated (Autor 2015). These differences between AI and prior technologies point to the differences in how humans can respond to the introduction of prior technologies and to the introduction of AI systems. Since prior generations of technologies are superior to humans in routine manual and standardized mental tasks, humans have found it increasingly difficult to compete with these technologies, and have been replaced by these technologies for those tasks and have had to switch to other jobs/tasks (Akst 2013, Brynjolfsson and McAfee 2014, Autor 2015, Acemoglu and Restrepo 2018). In contrast, when AI systems are introduced, especially in creative tasks, humans have an option to quit, but they can also compete with AI systems by leveraging their creativity, intuition, imagination and emotional expression (Brynjolfsson and McAfee 2012, Eberhard et al. 2017, Hertzmann 2018, Huang et al. 2019, Lu et al. 2020). This study contributes to this research stream by examining how different

human designers respond to the introduction of the AI system for a creative task on a crowdsourcing platform.

Prior research on AI limitations indicates that while modern AI systems with advanced deep learning capabilities are very impressive, humans are still more advanced in such qualities as creativity, imagination and emotions in general (Braga and Logan 2017), and creativity and emotional and social intentions in design and art specifically (Hertzmann 2018, Mazzone and Elgammal 2019). Recent advances in Generative Adversarial Networks (i.e., so-called Creative Adversarial Networks) suggest that algorithms can be trained to use the same distribution of styles used by human artists, but at the same time to maximize the differences between a new algorithmically generated art and all prior works, thus, making the AI-generated art as novel as possible (Elgammal et al. 2017). However, there are profound differences between machine “creativity” and human creativity. Mazzone and Elgammal (2019) highlight that a machine uses a combination of given elements as training sets without an outside reference, while a human artist gets inspiration from something in the outside world (e.g., nature). Additionally, Hertzmann (2018) argues that an important role of the artist is to supply the “intent” and the “idea” for the work. Hertzmann (2018) also notes that human artists possess creativity, growth, and responsiveness, and to achieve human level of art creation an AI machine needs to have capacity for consciousness, emotions, and social relationships.

Despite the current limitations of AI systems for design, they are increasingly being deployed by a variety of platforms for creative tasks. This study contributes to this stream of research on AI capabilities and limitations by focusing on the specific AI

system for logo design. While the primary goal of this study is to shed light on how human designers respond to the introduction of an AI system for logo design, I also seek to understand how the logos designed by human designers are different from AI-generated logos. Understanding the differences between AI-generated logos and human logos will also shed light on specific skills that human designers need to develop to successfully respond to the introduction of AI systems.

4.2.4 Emotions and Complexity

Finally, my research draws upon seminal work in the area of art perception to understand how designers might respond to the AI system. Berlyne's psychobiological model of aesthetic experience in art perception (Berlyne 1974, Marin et al. 2016) provides useful constructs relevant to my research context. Specifically, Berlyne's model proposes two interrelated constructs that affect human aesthetic experience in art – complexity and emotions (specifically, arousal or excitement). More recent research that uses Berlyne's model (see, for example, Marin et al. 2016) finds that complexity and arousal are positively associated in all conditions. The authors (Marin et al. 2016) also find that complexity is positively associated with beauty. More recent research in marketing finds that higher design complexity of an ad image helps increase attention to both the pictorial and the advertisement as a whole, positively affects ad comprehensibility, and attitude towards the ad (Pieters et al. 2010). More specific to logo design, Grinsven and Das (2016) find that logo complexity positively affects long-term brand recognition and brand attitude in the case of repeated exposures. Researchers (Salgado-Montejo et al. 2014, Bajaj and Bond 2018, Kim and Lim 2019) also find that positive emotions expressed in logos have positive effects on the attitude

towards brands (De Marchis et al. 2018). De Marchis et al. (2018) find that excitement and happiness are strongly associated with logo aesthetic attraction. Additionally, prior research finds that there is positive correlation between subjective human-rated logo complexity and logo “emotionality” (De Marchis et al. 2018) and positive relationship between logo complexity and excitement (Bajaj and Bond 2018).

This study continues this line of research and seeks to understand whether emotions and complexity of design are important variables in the context of crowdsourcing contests and seeks to understand how contest participants can leverage those attributes in response to the introduction of the AI system.

4.3 Research Context and Data

The crowdsourcing platform for design tasks allows clients to create design contests and allows designers to submit solutions for a monetary award. Logo-design contests constitute the main category of contests (90% of contests), while other categories (10%) include contests mostly for design of T-shirts. The basic lower-tier contests have an award amount below 110 US dollars and mostly seek simple design solutions. Clients with more-complex requirements typically choose a higher award amount (from 110 US dollars up to 2,284 US dollars, see Table 4.1). Interestingly, at the beginning of April 2018 the focal platform introduced an Artificial Intelligence logo maker that offers hundreds of logo design solutions based on a client’s inputs such as a company name, a slogan, preferred styles, colors, and shapes. The whole process of logo generation using the AI system takes about 5 minutes. Once a client has made a choice, he or she can purchase the logo created by the AI system and acquire all the rights to use the logo. The basic price of an AI-generated logo is 20 US dollars, while the full

resolution logo with different varieties of format costs 65 US dollars. This pricing is close to the pricing of lower-tier contests where clients with simple requirements can invite solutions from designers on the platform. The whole process for a contest with human designers may take up to 10 days to complete. There are advantages and disadvantages of using human designers versus using the AI logo maker. On the one hand, the AI logo maker is very fast, and as described by its creators, it constantly learns as more people use it and it follows the most recent trends in logo design. On the other hand, a contest with human designers might provide more suitable logos with potentially better quality from professional designers. This setting provides a unique opportunity to understand the impact of the AI system on individual users' behaviors and strategies in response to the AI system introduction.

I collect data for the period of six months before the introduction of the AI logo system (from October 2017 till March 2018), and for the period of six months after the introduction of the AI logo system (from April 2018 to September 2018). As noted earlier, the AI system was launched at the beginning of April 2018. The dataset includes all 5,737 contests for logo design and other types of design contests for that period and the complete history of participation for 9,280 designers, which includes a total of 425,475 design submissions. To initiate a contest, a client (contest holder) needs to provide an award amount, a task description that includes a name to use in a logo, description of target audience, organization or a product and any specific requirements. It should be noted that 98.3% of contests are “hidden” which means that contest participants (designers) cannot see design submissions by other designers until the end of a contest.

During a contest, designers submit their solutions, and a contest holder provides star ratings to select submissions. Since most contests are “hidden”, participants can only see other designers’ profiles (and experience), order of their submissions and star ratings for submissions but not the submissions themselves. At the end of a contest, a contest holder announces a winning submission.

Variables’ definitions and descriptive statistics are shown in Table 4.1.

4.4 Methodology

First, I compare the overall number of contests of each type available on the platform in both periods. Next, I compare contests’ task descriptions before and after the AI launch to understand which contests are more likely to be replaced or “cannibalized” by the AI system.

4.4.1 Task Descriptions’ Requirements Classification

To begin with, I manually classify task requirements in 300 random logo contest descriptions. The main categories that I observe are: concrete or specific logo requirements, such as “I want the picture of a child with the graduation cap as the shadow”; abstract logo requirements, such as “Overall design must be sleek and classy”; requirements to convey brand emotions/feel, such as “I’d like the logo to convey happiness and excitement”. I validate these categories using the Amazon SageMaker tool¹⁹ (each contest description was classified by three users) and find that the percentage of agreement between the manual classification and Amazon

¹⁹ <https://aws.amazon.com/sagemaker/>

SageMaker users' classification is 90%. To automatically extract those categories in task descriptions for all 5,737 contest descriptions, I use the following methods.

First, to capture brand emotions/feel requirements, I use sentiment analysis. I use the VADER (Valence Aware Dictionary and sEntiment Reasoner) sentiment toolkit (Hutto and Gilbert 2014) implemented in Python²⁰. Hutto and Gilbert (2014) have shown that the VADER sentiment analysis is very accurate and outperforms human raters. I expect that contest descriptions with more requirements to convey brand emotions/feel will have a higher overall sentiment score as compared to contest descriptions with fewer requirements to convey brand emotions/feel. I confirm this method by comparing task descriptions and sentiment scores for a sample of manually classified 300 random contest descriptions.

To capture abstract or specific (concrete) requirements, I use the largest available dataset of 10,000 abstract/concrete terms (Pexman et al. 2017). The proportion of abstract/concrete words in that dataset is very close to 50/50. Concrete and abstract words in the dataset are nouns, verbs, adjectives, and adverbs. For instance, such words as “bottle, mug, spade” are concrete, while such words as “excitement, happiness, truth, trust, hope” are abstract. I count the number of abstract and concrete (specific) terms (from the list of 10,000 terms) and use those word counts as a proxy of abstract and specific requirements in each contest task description. Again, I confirm this method by comparing task descriptions and the count of abstract and specific terms for a sample of manually classified 300 random contest descriptions.

²⁰ <https://github.com/cjhutto/vaderSentiment>

4.4.2 Complexity of Design Images

Pieters et al. (2010) and Grinsven and Das (2016) use human raters to measure perceptions of ads' complexity and logo complexity. Since I have 425,475 design images in my dataset, it would be infeasible to manually classify all these images as more-complex or less-complex. Hence, I employ a measure widely used in computer science literature to measure image complexity - a measure termed a spatial information (SI) complexity (Yu and Winkler 2013). Specifically, I use a mean of edge magnitudes from edges extracted using horizontal and vertical Sobel filters. The mean of edge magnitudes has shown better performance than other measures (Yu and Winkler 2013, Athar and Wang 2019)²¹. To validate this measure for the purpose of this study, I use the Amazon SageMaker tool (three Amazon SageMaker users classify each design) to classify 100 random designs using a Likert 1-5 scale (for instance, see Pieters et al. 2010) with the anchor points ranging from "Very simple" to "Very complex" and check the correlation of the automated spatial information complexity measure SI with the classification scores by the SageMaker users.

I find that the automated measure of complexity has high correlation (0.8) with the measure obtained by using SageMaker users. It is pertinent to note that the agreement on complexity among SageMaker users is 0.84, indicating that the automated measure is very close to the users' perception of complexity in terms of accuracy. Hence, in the analyses I use the automated measure of complexity calculated for 425,475 design images in the dataset (see examples in Figure 4.1).

²¹ The SI measure does not consider color of images, Ciocca et al. (2015) have shown that color does not influence the human perception of image complexity.

4.4.3 Emotions in Design Images

To measure emotions in design submissions, I use the following methods. First, I ask Amazon SageMaker users to evaluate 2,100 design images on whether they feel any emotions by looking at the image or do not feel any emotions. Each design is evaluated by three Amazon SageMaker users. Out of 2,100 designs, 1,036 are evaluated as “not eliciting any emotions”. Next, I use a deep learning model (using “adjective-noun pairs” features described below) that would predict each design as “eliciting emotions or not eliciting emotions”. After eliminating designs that do not have sufficient discriminatory power, the final training set includes 1,243 designs with 455 design images that belong to a category “not eliciting emotions”. Examples of design images with emotional content and without emotional content are shown on Figure 4.2. The model accuracy with this training sample is 78.7% (with a balance for precision and recall), which is close to accuracy (agreement) of 80% among SageMaker users. Thus, I use this model to predict whether a design image has emotional content or does not have emotional content. I perform predictions for all design images in the sample.

A second measure of design emotions is more granular. I build a deep learning model for 5 emotions based on the largest database of images classified in a prior study by human raters into the following emotions: amusement, awe, contentment, excitement, sadness (You et al. 2016)²². Since this dataset is unbalanced with some emotions being dominant, I use a more balanced dataset (4,865 images) that has

²² You et al. (2016) classify a total of 8 emotions - amusement, anger, awe, contentment, disgust, excitement, fear, sadness. However, Salgado-Montejo et al. (2014) find that emotions such as “fear”, “anger” and “disgust” are unlikely to be present in logos. I verify that this is the case in my context as well and focus on the 5 emotions.

approximately the same number of images for each emotion category. Additionally, prior research indicates that there is an “affective” gap between low-level features of an image (such as colors) and emotions that humans perceive when they look at an image. To address the “affective gap”, Borth et al. (2013) propose mid-level representations of an image based on adjective-noun pairs (ANPs). The authors use a deep learning model (SentiBank²³) to extract 1,200 adjective noun-pairs such as “colorful lights”, “great adventure”, “pleasant surprise” from each image and show that those adjective-noun pairs could be used for classifying the same set of emotions that I focus on in this study. In keeping with this, I use the SentiBank deep learning model to extract 1,200 adjective-noun pairs and their probabilities for each image and retain top 10 ANPs²⁴ for each image to use those as features in another neural network model that classifies each design image into one of the 5 emotions. It is pertinent to note that the task of automatic emotion extraction from images is an ongoing research area and that I reach classification accuracy levels close to the state-of-the-art methods (Yang et al. 2017, He and Zhang 2018, Liu et al. 2019).

Table 4.2 shows the confusion matrix (the trade-off between true positive rates, true negative rates, false positive rates, and false negative rates) for the model with 5 emotions with the total accuracy of classification 72.14%, which is much better than a random guess of 20%.

It is pertinent to note that in the prior task of labeling design images by SageMaker users (for presence or absence of emotions in design images), agreement among

²³ <http://www.ee.columbia.edu/ln/dvmm/vso/download/sentibank.html>

²⁴ Using top 10 Adjective-Noun Pairs (as opposed to top 15 or top 20, top 30, and so on) helps achieve higher accuracy of classification.

SageMaker users was close to 80%. Thus, the accuracy of 72.14% is comparable to SageMaker users' accuracy. Additionally, if the model predicts low probability for each of the 5 emotions, I assign it to the category "other/neutral" and control for this category in subsequent econometric analyses.

Table 4.3 shows descriptive statistics for the emotions and complexity variables.

For a comparison between emotions and complexity of human logo designs and emotions and complexity of AI-generated logo designs see Appendix – Section A5.

4.4.4 Relationship between Emotions and Complexity

Prior research has found that there is positive correlation between logo image emotions and logo image complexity (De Marchis et al. 2018). I confirm that positive correlation in this setting. I find that the correlation between a binary variable for presence/absence of emotions and complexity is 0.32 for both higher-tier and lower-tier logo contests and 0.2762 for non-logo contests, which is between the reported values in the prior study (De Marchis et al. 2018) that range from 0.16 (for objective non-human-rated measures of complexity and subjective measures of emotions) to 0.54 (for subjective human-rated measures of complexity and subjective measures of emotions).

4.5 Results

4.5.1 Effects of AI on the Number and Composition of Contests

First, I compare the number of lower-tier logo contests (with the award below the median value of 110 US dollars), the number of higher-tier logo contests (with the

award at or above 110 US dollars) and the number of non-logo contests before and after the AI launch.

Results indicate that after the introduction of the AI system for logo design the number of lower-tier contests for logo design (with human designers) decreases by 25%, while the number of higher-tier contests for logo design decreases by 5%. In contrast, the number of non-logo contests increases by 10%. Figure 4.3 shows the comparison of lower-tier and higher-tier logo contests before and after the AI launch.

Next, I compare task descriptions of lower-tier and higher-tier logo contests before and after the AI launch to see if distributions of categories in task requirements change after the AI launch. Table 4.4 shows the comparison of lower-tier and higher-tier logo contest task descriptions on the dimensions of “abstract”, “specific” task requirements and requirements to convey “brand emotions/feel”.

Table 4.4 shows that higher-tier logo contests have more abstract requirements (36.9% more before AI and 27.6% more after AI) and more requirements to convey brand emotions/feel (9.56% more before AI and 6.3% more after AI) as compared to lower-tier logo contests. Also, the average number of abstract requirements per lower-tier contest increases from 6.5 to 7.07 after the AI launch, a 9% increase. This suggests that contests with fewer abstract requirements are more likely to be “cannibalized” by the AI system. The results in Table 4.4 are confirmed in regression analyses.

The number of “brand emotions/feel” requirements (sentiment score) per lower-tier contest increases from 0.659 to 0.692 (5% increase). Overall, since the number of lower-tier logo contests reduces by 25% after the AI launch and since the “cannibalized” contests have fewer abstract requirements and fewer requirements to

convey brand emotions/feel as compared to the remaining contests (lower- and higher-tier contests), any designer has to choose among more-complex contests to participate in after the AI launch.

4.5.2 Designers' Responses to the Launch of the AI system for Logo Design

I seek to understand the behaviors of designers relating to their choice of contests to participate in, after the introduction of the AI system. First, I check the attrition rate in lower-tier logo contests, higher-tier logo contests and non-logo contests before and after the introduction of the AI system. I calculate the number of distinct designers participating in contests 6 months before the AI launch and compare that to the number of the same designers still participating in contests 1 month before the AI launch. Similarly, I calculate the number of distinct designers participating in contests 1 month after the AI launch and compare that to the number of the same designers still participating in contests 6 months after the AI launch. I find that after the AI launch the attrition rate of designers increases from 65% to 69.7% (the difference is 7.23%) in the lower-tier logo contests, decreases from 59% to 57.45% (the difference is 2.6%) in the higher-tier logo contests, and increases marginally from 79% to 80% (the difference is 1.26%) in the non-logo contests.

To examine the designers' behaviors, I track 2,374 designers who are active in both periods – before and after the AI launch. I find that 33.5% of designers switch to higher-tier logo contests (I term these cross-tier designers), while other designers (18%) switch to other non-logo contests (I term these cross-category designers). I define “switching” as an increase in the proportion of contests (higher-tier logo contests or non-logo contests) in which each designer participates in after the AI introduction. I find that

there are three major *groups* of designers after the AI launch: designers who continue to participate mostly in lower-tier logo contests (I term these focused designers), designers who switch vertically to higher-tier logo contests (i.e., cross-tier designers) and designers who switch horizontally to non-logo contests (i.e., cross-category designers). Only 3% of designers switch both vertically and horizontally. I report the results separately for each group and find the results to be consistent when I consider only cross-tier or only cross-category designers and exclude those 3% of designers who switch in both directions. Additionally, there is a group of designers (about 10% of designers, or 240 designers who made 6,451 design submissions in my dataset) who participated only in non-logo contests both before and after the introduction of the AI system. I use the designers in that group as a control group in the difference-in-differences models with propensity score matching (Smith and Todd 2005; Liu and Lynch 2011).

To examine the switching behavior of designers in response to the introduction of the AI system, I match designers on their experience and their activity (number of submissions) on the platform using propensity score matching and dichotomize proportions of higher-tier logo contests and non-logo contests (before the AI launch) by median splits. The results are reported in Table 4.5. I find that cross-tier designers who had prior exposure to higher-tier contests before the AI launch are more likely to switch to higher-tier logo contests, while cross-category designers who had prior exposure to non-logo contests prior to the AI launch are more likely to switch to non-logo contests upon the introduction of the AI system, as compared to the focused designers who are more likely to continue to participate in lower-tier logo contests.

4.5.3 Successful Designers' Responses to the Introduction of the AI System for Logo Design

Next, I seek to understand how the behaviors of successful designers differ from the behaviors of the other designers after the introduction of the AI system. It is pertinent to note that the successful designers include designers who were winning before the introduction of the AI system (and keep winning after the AI launch) as well as new winners, while all other designers are categorized as unsuccessful (unsuccessful after the AI launch).

First, I want to understand whether participation in contests and the number of submissions change for the three groups of designers and whether successful designers are different in their behaviors from the unsuccessful designers. The general model has the following form:

$$Y = \text{After} + \alpha_u \quad (4.1)$$

where Y – dependent variable (first set of variables of interest - the number of contests per user per day, number of submissions per user per contest; second set - emotions and complexity of designs); After – a dummy variable indicating a period after the introduction of the AI system; α_u – designer fixed effects.

Table 4.6 and Table 4.7 show comparison of the successful and the unsuccessful designers in the three groups before and after the AI launch in models with designer fixed effects. The omitted group is the unsuccessful designers in the period before the AI launch.

The results (Table 4.6 and Table 4.7) show that the unsuccessful designers participate in more contests after the AI launch (increases by 15% for focused designers, increases by 14% for cross-tier designers, and increases by 13.3% for cross-

category designers), and either decrease or do not change the number of submissions per user per contest (decreases by 31% for focused designers, does not change for cross-tier designers, and does not change for cross-category designers). In contrast, the successful designers (comparing “successful designers before AI” and “successful designers after AI” in the tables) almost do not change the number of contests that they participate in after the AI launch (all changes are less than 1% for the three groups of designers). However, they substantially increase the number of design re-submissions per contest (increases by 60% for focused designers, increases by 30% for cross-tier designers, and increases by 55.7% for cross-category designers).

Next, I seek to understand whether successful designers and unsuccessful designers increase emotional content and complexity of design submissions in response to the introduction of the AI system²⁵.

Interestingly, I find that the successful designers increase emotional content (Table 4.8) and the complexity of their design submissions (Table 4.9) after the introduction of the AI system. These results are consistent across all three categories of designers.

As for the effect sizes (Table 4.8 and Table 4.9), successful focused designers increase emotional content by 6.5% and complexity by 7.5% after the AI launch as compared to the period before the AI launch. Successful cross-tier designers increase emotional content by 9.3% and complexity by 11.57% after the AI launch as compared

²⁵ Prior to this, I examine if the emotional content and complexity of design submissions do indeed have a significant influence on the probability of winning a contest. As shown in the Appendix, Section A6, I find that both the emotional content as well as the complexity of the design submissions have a significant impact on the likelihood of winning a contest.

to the period before the AI launch, and successful cross-category designers increase emotional content by 14.37% and complexity by 17.1% after the AI launch.

The logit model with five emotions (Table 4.10) shows that all three groups of successful designers increase the “excitement” of their design submissions after the AI launch.

Additionally, I find that the unsuccessful designers in all three categories do not change emotional content of their designs (Table 4.11) and do not change complexity of their design submissions (Table 4.12) after the AI launch.

The granular model with five emotions (Table 4.13) also shows that the unsuccessful designers in all three categories do not change emotional content of their design submissions.

4.5.4 Additional Controls

I add several additional controls that could affect the observed outcomes. For example, the AI system might have changed competition in contests, or changed contests’ task requirements due to partial cannibalization of lower-tier logo contests. Additionally, some designers change their contest choices after the AI introduction by switching to higher-tier logo contests or to contests in other non-logo categories. Those changes might be responsible for some of the results reported in the prior section, so I control for those changes to support the main results.

Control for Changes in Competition. The introduction of the AI system might have changed competitive dynamics in contests. If competition in contests changes after the AI introduction, designers might respond to changes in competition, rather than to the

introduction of the AI system. To account for this, I first check whether increased competition has an impact on the dependent variables of interest such as emotional content and complexity of designs. I find that when the total number of designers and the number of total design submissions increases in a contest, designers do not significantly change emotional content and complexity of their designs in that contest. This result is supported for both successful and unsuccessful designers - both before the AI introduction and after the AI introduction (for example, for the effects of the number of designers and the number of submissions per contest, see Appendix, Table A8). Nevertheless, I still control for changes in competition. To do so, I use propensity score matching (Rosenbaum and Rubin 1983) to match contests before and after the AI introduction on such variables as the number of designers in each contest and the number of design submissions in each contest.

Control for Changes in Task Requirements. If some of the task requirements change in contests after the introduction of the AI system, designers might change their behaviors because of the changes in the task requirements. In examining the impact of task requirements' changes, I find that if task requirements increase, the dependent variables of interest such as emotions and complexity increase by a small amount (i.e., the coefficients have small effect sizes, see Appendix, Table A9). To control for potential changes in task requirements, I match contests (using propensity score matching) before and after the AI launch on these 3 classes of requirements - specific requirements, abstract requirements, and requirements to convey brand emotions/feel.

Control for Choice of Contests (Switching). In addition, cross-tier designers and cross-category designers might change their strategies because they switch to contests

with more-complex requirements after the AI system introduction. To address this issue, I compare only matched logo contests in the higher-tier before and after the AI launch for cross-tier designers and compare matched logo contests (i.e., exclude non-logo contests) before and after the AI launch for cross-category designers.

The results for the “before-after” analyses with propensity score matching and with the controls for competition, task requirements and switching are reported in Tables 4.6, 4.7, 4.8, 4.9, 4.10, 4.11, 4.12, and 4.13 (in columns with PSM designation). They are consistent with the results for the earlier “before-after” analyses that were reported before the addition of these controls.

4.5.5 Difference-in-differences Analysis with Propensity Score Matching

To support the causal identification of the effects, I use a difference-in-differences (DiD) model with propensity score matching (PSM) techniques (Smith and Todd 2005, Liu and Lynch 2011). The general formula for the model is the following:

$$Y = \text{PSM} (\text{After} + \text{Treatment_group} + \text{After} * \text{Treatment_group}) \quad (4.2)$$

where Y – dependent variable (the number of contests per user per day, number of submissions per user per contest, emotions, and complexity of designs); PSM – propensity score matching; After – a dummy variable indicating a period after the introduction of the AI system; Treatment_group – a group of designers (focused designers, cross-tier designers, and cross-category designers) that were affected by the introduction of the AI system for logo design.

The control group in this case includes designers who always participated in non-logo contests both before and after the introduction of the AI system for logo design. Since the AI system is designed for logo design only, designers participating in non-

logo design contests should not be affected by the introduction of the AI system for logo design. For the matched control group, I restrict the sample to only successful designers when I perform analyses for the responses of successful designers (in the treatment groups affected by the AI introduction). Similarly, I restrict the sample to only unsuccessful designers in both treatment and control groups when I perform analyses for the responses of unsuccessful designers to the AI introduction. First, consistent with the prior section, I match contests before and after the AI launch on such variables as task abstract/specific requirements and requirements to convey brand emotions/feel, the number of designers and the number of submissions per contest. To control for contest switching, I match only higher-tier contests before and after the AI introduction when I perform analyses for cross-tier designers. Similarly, I match only logo contests before and after the AI introduction when analyzing cross-category designers. Next, I match users in treatment and control groups on their experience with the platform using the variable *Experience*, which represents the number of days since registration for each designer (see Table 4.1). The purpose of the matching is to compare treatment and control groups in very similar contest conditions and for very similar designers, to isolate the effect of the AI introduction on designers' changes in behaviors and strategies.

After the matching I perform a difference-in-differences analysis on the matched groups. An important test for the difference-in-differences method is a test of parallel trends before the treatment (Autor 2003, Ryan et al. 2015). I perform the tests for each group (focused designers, cross-tier designers and cross-category designers) by interacting the treatment group with the monthly dummies in the period before the

introduction of the AI system (periods from t-6 to t-1) and find that the slopes for the treatment and control groups are not significantly different, which confirms the parallel trends assumption (see tables A10, A11 and A12 in the Appendix).

I use difference-in-differences methods to confirm prior results from Table 4.6 and Table 4.7 regarding the number of contests that each designer participates in per day before and after the AI launch, and the number of re-submissions per designer per contest before and after the AI introduction. All the results are confirmed and are reported in the Appendix (Table A13 and Table A14).

I also confirm that the successful designers in each group increase the emotional content (Appendix, Table A15) as well as the complexity of their design submissions (Appendix, Table A16). It should be noted that I report the results for the difference-in-differences models without user or time fixed effects but adding the user and time (month) fixed effects provides consistent estimates.

It is pertinent to note that in the difference-in-differences models with propensity score matching (Table A15 and Table A16 in Appendix) the effect sizes are larger in almost every case as compared to the “before-after” models with designer fixed effects (Table 4.8 and Table 4.9).

I confirm the results for *Excitement* (for the successful designers) in difference-in-differences models (Appendix, Table A17). For other emotions, the results for PSM-DID are not significant (consistent with Table 4.10).

The difference-in-differences models with propensity score matching also confirm that the unsuccessful designers in all three groups do not change the emotional content or the complexity of their design submissions after the AI introduction.

4.5.6 Robustness Checks

I perform additional robustness checks to support the main results. For example, the emotions and complexity of designs might have increased because there is a time trend, and the effect is still not captured in the difference-in-differences models. To address this issue, I compare short periods of 2 months before and 2 months after the AI system launch and find that the results are highly consistent with the longer periods of 6 months before and 6 months after the AI launch. As a falsification test, I compare the periods of 3-4 months (and 5-6 months) before the AI launch with the period of 1-2 months before the AI launch. In the latter case, since the AI was not introduced at that time, there should be no change in emotions and complexity of designs even among the successful designers, and I confirm that this is the case. All the robustness checks confirm the main findings – the successful designers consistently increase emotional content and complexity of their designs after the introduction of the AI system.

Since I find that the successful designers increase the number of re-submissions per contest after the AI launch, I also check whether the number of re-submissions is associated with the improvement in quality for the successful and the unsuccessful designers after the AI launch. I find that the successful designers increase emotional content of their designs with each additional submission (increases by 2.02% for each re-submission in the focused designers' group, increases by 2.27% for each re-submission in the cross-tier designers' group and increases by 2.33% for each re-submission in the cross-category designers' group) and increase complexity of their designs with each additional submission (increases by 2.15% for each re-submission in the focused designers' group, increases by 2.42% for each re-submission in the cross-

tier designers' group and increases by 2.97% for each re-submission in the cross-category designers' group). In contrast, I find that the unsuccessful designers do not increase emotions and complexity with each additional re-submission in each group of designers.

An alternative explanation for the increase in the number of re-submissions per contest by the successful designers as a response to the AI system could be the differences in feedback patterns after the AI launch as compared to the period before the AI launch. First, I check and find that if a client (i.e., a contest holder) provides a high star (4 or 5 star) rating to select submissions, then those designers who got the high star rating have a higher likelihood of making a re-submission in the same contest as compared to the designers who did not receive a high star rating. Next, I compare high-star feedback frequency before and after the AI launch, and find that the feedback frequency changes only marginally after the AI launch. Finally, as a robustness check, I also match contests before and after the AI launch on the number of high star ratings provided by contest holders to each designer and find the results to be consistent with my findings.

4.6 Implications and Conclusion

This study seeks to understand the overall impact of the AI logo system on the contests available on the crowdsourcing platform and on designers' strategies in response to the AI system launch. As noted earlier, AI systems are different from prior generations of technologies in their impact on humans' behaviors and strategies. Given that prior technologies are superior to humans in routine and standardized tasks, the only option for humans is to quit those tasks or switch to other tasks. However, AI systems simulate

non-routine tasks when “tacit” human knowledge is required. Humans can quit those tasks, or they can compete with the AI systems by leveraging their creativity, imagination, emotional expression, empathy etc.

I find that the AI system “cannibalizes” lower-tier logo contests with less-complex requirements to a greater extent as compared to higher-tier logo contests with more-complex requirements. This effect increases the overall complexity of remaining logo contests available on the platform. I find that different groups of designers respond differently to the introduction of the AI system. The focused designers in the first group either leave the platform or continue to participate predominantly in lower-tier logo contests, while the cross-tier designers and cross-category designers switch to higher-tier logo contests or to non-logo contests accordingly. I find that the designers who had prior exposure to higher-tier logo contests or non-logo contests are more likely to switch to higher-tier or to non-logo contests accordingly. I use “before-after” analyses with designer fixed effects and employ propensity score matching as well as difference-in-differences models to support causal identification of the effects. Interestingly, I find that the successful designers become more focused and increase the number of re-submissions per contest substantially (by 30-60% depending on a group of designers) as compared to the unsuccessful designers after the AI launch. In contrast, the unsuccessful designers participate in more contests after the AI launch (by 13-15%) and either decrease the number of re-submissions per contest (in the focused designers’ group) or do not change the number of re-submissions per contest (in the groups of cross-tier designers and cross-category designers). Further, I find that the successful designers, as compared to the unsuccessful designers, in all three groups increase

emotional content (by 6.5-14.37% depending on a group of designers) and complexity (by 7.5-17.1% depending on a group of designers) of design re-submissions as a response to the introduction of the AI system. Finally, I find that the introduced AI system's designs are different from human designs in several ways (see Appendix – Section A5). With respect to emotions, humans can produce logos with more positive emotions and fewer neutral emotions as well as generate emotions with higher intensity as compared to the AI system. The AI system also has an upper limit of complexity, and about half of the designers on the platform can produce designs that have higher complexity compared to the complexity of the most complex designs of the AI system.

This study has important implications for theory and practice. First, prior research on the effects of AI adoption on employment has mostly been conducted at the industry or geographic region level, and the effects on employment are found to be negative. More recent studies at the firm level (for example, Dixon et al. 2019) point to more nuanced effects, when AI systems increase overall employment, but negatively affect employment of managers. This study contributes to this research stream by expanding the context of AI to crowdsourcing platforms. To the best of my knowledge, this study is among the first to explore users' successful strategies and responses to the introduction of an AI system in a decentralized crowdsourcing platform for design tasks.

Second, prior research has mostly used surveys to understand responses of employees to the adoption of AI systems at their workplace. The findings from prior studies are context specific. Some employees, for instance, disability care workers feel less threat from smart AI robots because they think that these AI robots are not

advanced in emulating human touch and emotions (Wolbring and Yumakulov 2014). In contrast, in the hotel industry employees feel more job insecurity in response to AI systems' adoption and are more likely to quit (Li et al. 2019). This study extends this line of research into empirical setting. The results show that, in the context of a crowdsourcing platform for design tasks, successful designers (freelancers) respond to the AI system by focusing on each contest (i.e., making more re-submissions per contest) and by increasing emotional content and complexity of their design re-submissions.

Next, the findings of this study contribute to the research related to design emotions and complexity. Prior research has explored these variables in laboratory and experimental settings (Grinsven and Das 2016, Bajaj and Bond 2018, De Marchis et al. 2018). Additionally, Berlyne's model of aesthetic perception of art considers emotions and complexity as the two key variables. This stream of research has found positive correlation between emotions and complexity. To the best of my knowledge, this study is among the first to measure design emotions and complexity empirically on a large scale in a decentralized platform for design crowdsourcing. I find that an increase in emotional content and in complexity of designs increases the probability of winning a contest (see Appendix – Section A6). I also show that there is positive correlation between logo emotional content and complexity. Thus, this study empirically confirms prior findings in the new context of a crowdsourcing platform for design tasks and contributes to that research stream.

The findings of this study also have valuable practical implications. Managers can use the findings to get insights into strategic behaviors of users on crowdsourcing

platforms in response to the adoption of an AI system. Additionally, the findings can help managers to better segment clients into those who would prefer the AI system and clients who should use human designers for more-complex logos with more emotional content. Importantly, market providers can use my findings to optimize revenue from both sources – AI and human designers by designing flexible pricing and marketing mechanisms. Another source of revenue could be a hybrid model when a client starts using the AI system, and then an expert designer finalizes the design. Additionally, in the presence of the AI system managers might offer new incentives to top-performing designers to reduce turnover rates. Finally, knowledge of the capabilities and limitations of AI systems for design might be helpful for human designers who can leverage such design attributes as emotions and complexity to improve their outcomes.

This study is not without limitations. First, I analyze the data from a single crowdsourcing platform for design tasks. Future research might look at other platforms and compare the results with my findings. Second, I measure emotions and complexity of a whole image. Future research might investigate which specific parts of an image make it more “emotional” or more complex and how those parts might be related to each other and to the probability of winning a contest. Finally, although I observe responses of designers to the AI system in the period of 6 months after the AI launch, future research might explore which strategies (such as switching to participate in more-complex tasks in the same category or across categories) are better in the longer term.

Chapter 5: Conclusion

The last decade has seen a tremendous growth in the number of online platforms. Commensurate with the growth of these platforms, there has been an increasing interest among academicians and practitioners to study these platforms from multiple perspectives and in many disciplines.

Prior research on online platforms has considered various aspects of those platforms such as their design and architecture, interactions among users and users' networks, network externalities, impact of platforms on society and platform governance (Constantinides et al. 2018, de Reuver et al. 2018). However, less attention has been devoted to users' strategies and relationships between strategies and outcomes. My dissertation seeks to fill this gap by focusing on understanding users' successful strategies and behaviors and the impacts of these on outcomes on social media and crowdsourcing platforms. Leveraging large-scale data and state-of-the-art techniques for analyzing unstructured data helps me understand users' successful strategies in such platforms on a more granular level.

Specifically, social media research has explored extensively interactions among users (Kietzmann et al. 2011), but less attention has been devoted to how firms use social media. And there are even fewer studies that focus on competing firms' strategies on social media and on the effects that those strategies have on related outcomes. My first essay seeks to address this gap by focusing on content strategies of traditional rival firms on Twitter and by exploring which content strategies are more successful for those rival firms.

Similarly, research on crowdsourcing platforms and corresponding contests on those platforms has studied design of those contests such as prize structure, types of contests and feedback mechanism, but there are hardly any studies that focus on users' successful strategies on these crowdsourcing platforms. The second and third essay in this dissertation seek to fill the gap by exploring users' strategies on crowdsourcing platforms for design tasks and by connecting those strategies with the likelihood of users' success. Specifically, the second essay focuses on strategies of experienced and less-experienced users in leveraging information spillovers in open contests on a crowdsourcing platform. And the third essay focuses on understanding users' successful strategies in response to the introduction of an AI system on a crowdsourcing platform.

Collectively, my findings indicate that there are a variety of successful strategies on these platforms. Importantly, those strategies are tied to certain design elements and unique functionalities of these platforms. For instance, the firms increase their online engagement and attract more new followers by leveraging specific capabilities of Twitter, such as using hashtags to create communities of users around specific topics and using offline-online campaigns and online contests to harness the power of "crowd" and to invite users to participate in co-creation and co-innovation activities. In crowdsourcing platforms such design solution as visibility of prior submissions creates interesting opportunities for recombinant innovations when a client can benefit from the collective "crowd" wisdom of designers participating in a contest. Interestingly, the success in such contests depends on the ability to leverage information spillovers and to recombine valuable information from several prior

submissions to create novel solutions that incorporate several valuable design elements that a client prefers. Finally, such a change of the crowdsourcing platform design as the introduction of an Artificial Intelligence system creates interesting changes in designers' behaviors who leverage their emotional expression and creativity to differentiate themselves from the competing AI system and, thus, increase their probability of success in contests.

Specifically, the first essay of this dissertation seeks to understand the successful strategies of traditional competing firms in engaging users in online brand communities. Traditional rival firms, by experiencing competitive pressures, exhibit Isomorphism, or similarity, in their supply and demand sides in offline channels. Recently, the research has shown that this Isomorphism is also present and could be observed among online Web footprints of competing firms. I extend this research to social media where firms might leverage unique affordances to engage with their followers and create brand communities with millions of followers. I specifically focus on both similarity and dissimilarity in rival firms' content strategies on Twitter and seek to understand whether being closer to or farther away from the content of close rivals is beneficial for a focal firm. I find that it is dissimilarity of a focal firm's content strategies with the content strategies of close rival firms, that provides a focal firm an edge on Twitter and significantly affects online engagement and creates growth in the userbase of online brand communities. By exploring the underlying mechanism, I find that dissimilar rival firms differentiate themselves from closely competing firms by leveraging "community" and "co-creation" capabilities of Twitter and by doing so positively affect online engagement and attract more new followers. This study

contributes to research on social media affordances and specifically to research on how rival firms can leverage those affordances and differentiate themselves from competing firms. The findings of this study also provide specific recommendations to practitioners seeking to increase their social media presence and online engagement.

The second essay of my dissertation seeks to understand the successful strategies of designers in open crowdsourcing contests for design tasks. An “open” nature of contests creates a possibility that prior ideas will be imitated by other designers. Thus, information spillovers are possible in these contests. Depending on whether these spillovers are leveraged fully or partially, a client might benefit from them or might penalize designers for direct imitation. I find that the successful strategy in such contests involves leveraging information spillovers from several prior highly-rated designs by recombining those spillovers to create novel and useful design solutions rather than “blindly” imitating individual designs. Interestingly, experienced designers are more adept at synthesizing or recombining information spillovers, while the less-experienced designers tend to imitate individual prior highly-rated designs. This study contributes to the growing literature on information spillovers in online contexts, and to the growing research on recombinant innovations. Practitioners can use my findings to understand the strategies of successful designers in open contests that involve synthesizing or recombining design elements from several highly-rated designs, which are only possible in “open” contests where prior solutions are visible to other designers. Importantly, managers can design better mechanisms and incentives to compensate early movers for positive externalities that they create, especially given that those externalities are beneficial for the clients.

The third essay of my dissertation seeks to understand the behavioral changes of designers in response to the competing Artificial Intelligence system introduced on a crowdsourcing platform for design tasks. Specifically, I focus on strategies that are more successful in this context. Interestingly, I find heterogeneous responses to the competing AI system. First, the strategies of unsuccessful designers involve participation in more contests after the introduction of the AI system. Concurrently, these designers do not change the quality of their design submissions after the AI launch. Importantly, the successful designers adopt a completely different strategy by focusing on the elements of designs in which humans have the comparative advantage over the AI system – emotional expression and complexity. The successful designers increase emotional expression and complexity of designs as a response to the introduction of the AI system by focusing on each contest and making more re-submissions within each contest. This study is among the first to provide insights into the successful responses to the competing AI system by users in a decentralized marketplace, and, thus, it contributes to the nascent research focusing on the impact of AI on users' behaviors and strategies in this context. Importantly, the findings of this study provide valuable insights to platform managers who can use my findings to understand nuanced heterogeneous responses to AI by designers and to leverage these insights for improving the design of these platforms and for improving the co-existence of humans and AI systems with the goal of optimizing the revenue from both sources.

Collectively, the findings of the three essays of this dissertation make contributions to the research on online platforms, specifically, to the research focusing on the strategies of users that help them be more successful in those platforms. My studies

also have valuable practical implications. By leveraging granular large-scale data and state-of-the-art data analytics techniques, my studies uncover more nuanced strategies for users' success, which are beneficial for managers of these platforms and for industry practitioners.

Tables

Table 2.1. Descriptive statistics of the dataset

Sector	Number of firms	Total number of tweets	Average number of tweets per firm	Min. number of tweets for a firm	Max. number of tweets for a firm	Average (min:max) number of followers as of August 2017
Retail Trade	199	2,422,968 (893,525 firm-initiated tweets)	12,176 (4,278 firm-initiated tweets)	539 (365 firm-initiated tweets)	72,137 (17,723 firm-initiated tweets)	489,026 (1102:8,744,557)

Table 2.2. Descriptive statistics for the main variables

Variable name	Variable definition	Average value	Standard deviation
Favorites	Dependent variable. Represents the number of favorites (a.k.a likes) per firm per quarter	6,375.5	24,309.9
Retweets	Dependent variable. Represents the number of retweets per firm per quarter	3,088.7	8,811.3
New followers	Dependent variable. Represents the number of new followers per firm per quarter	10,584.4	23,697.5
Similarity	Independent variable. Represents cosine similarity of a focal firm's tweets with top competitors' tweets per quarter. The value is from 0 to 1, where 1 is the most similar	0.091	0.06
Tweets	Control variable. Represents the number of tweets per firm per quarter	193.7	167.6
Followers	Control variable. Represents the number of total followers per firm per quarter	151,489.5	299,779

Table 2.3. Fixed effects, GMM and 2SLS similarity model results for the period Jan. 2012-Aug. 2017

Model	FE	GMM	2SLS	FE	GMM	2SLS	FE	GMM	2SLS
Variables	Favorite	Favorite	Favorite	Retweet	Retweet	Retweet	New Follow	New Follow	New Follow
Similarity	-31475***	-53177***	-25394***	-11553***	-8399***	-9774**	-4302 ^{ns}	-13626*	-16488**
St. Er.	(8132)	(10673)	(9686)	(2871)	(888.75)	(3942)	(7256)	(7331)	(7420)
Tweets	39.5***	32.6***	35.99***	18.3***	18.74***	17.51***	9.45***	10.4**	8.02***
St. Er.	(2.65)	(8.62)	(4.6)	(0.935)	(2)	(1.6)	(2.36)	(4.7)	(2.09)
Followers	0.035***	0.0084**	0.034***	0.0079***	0.002***	0.0131***	-.003 ^{ns}	0.005 ^{ns}	0.052***
St. Er.	(0.002)	(0.0042)	(0.002)	(0.00073)	(0.00063)	(0.00092)	(0.002)	(0.004)	(0.002)
Sample	3509	3197	3353	3509	3197	3353	3509	3197	3353
Firms	156	156	156	156	156	156	156	156	156
R-squared	0.276		0.2789	0.329		0.35	0.0054		0.4451
*** $p < 0.01$ ** $p < 0.05$ * $p < 0.1$ ns – not signific.									

Note. Mean cosine similarity: 0.09, median: 0.08, standard deviation: 0.06, min: 0, max: 0.48; year and quarter fixed effects are not reported.

Table 2.4. Beta regression model with a logit link results for each higher-level category

Indep. var. \ Dep. var.	#Contests soliciting UGC	#Offline-online campaigns	#Special events	#Product collections	#Expert tips
Similarity	-1.774*** (0.354)	-2.14*** (0.211)	-1.73*** (0.419)	-1.02*** (0.382)	-1.249*** (0.08)
Constant	-2.08*** (0.016)	-2.08*** (0.016)	-2.08*** (0.016)	-2.08*** (0.016)	-2.08*** (0.016)

Note. Independent variable is percentage of tweets in each category. Dependent variable is similarity. The results for lower-level categories are mixed (2 categories, namely “events” and “questions”, have small negative coefficients while other 3 categories have positive coefficients)

Table 2.5. Normalized engagement (number of (favorites/retweets)/tweet/100,000 followers) by category

Tier	Tweet Category	Favorites	Retweets	Number of tweets
I	#Offline-online campaigns	19.74	10.07	13017
I	#Contests soliciting UGC	16.67	15.09	8773
II	#Expert tips	15.5	8.24	78334
II	#Product collections	14.28	9.24	7920
II	#Special events	13.3	8.01	8830
III	Product information	12.5	6.63	228207
III	Product usage tips	11.44	6.57	46503
III	Questions	11.02	7.31	96468
III	Events	10.74	6.82	22439
III	Coupons and promotions	9.92	8.69	130872
	Misc.	12.62	8.35	234289

Notes. The normalized engagement numbers are calculated after removing outliers in each category. All comparisons between categories are statistically significant.

Table 2.6. Two-stages least squares model results for the effect of usage of higher-level categories on online engagement and new followers' acquisition rate

Depend. Variables	#Contests soliciting UGC	#Offline-online campaigns	#Special events	#Prod. collections	#Expert tips
Favorites	4067.7** (1643)	1938.8** (812.2)	3541.9** (1597.7)	2784.9** (1127.5)	312.9** (130.2)
Retweets	1565.7** (657.6)	762.9** (323.5)	1393.7** (639.7)	1095.8** (460)	123.1** (52.5)
New Followers	7592.3*** (1988.3)	3817.6*** (947.5)	6974.4*** (1942.6)	5483.7*** (1299.7)	616.1*** (138.3)

Notes. Independent variables - percentage of tweets in each category. Favorites, retweets and new followers are dependent variables. *** $p < 0.01$; ** $p < 0.05$; * $p < 0.1$; ns – not significant.

Table 3.1. Descriptive statistics for the main variables

Name	Description	Mean	St. Dev.	Min	Max
Winner	A winner variable dummy indicating whether a submission by a designer is a winning submission in a contest				
Sim_insp	Similarity with all image inspirations in a contest. It represents averaged value for pairwise structural similarities of each submission in a contest with each image inspiration provided by a contest holder	0.523 (0.52)	0.2 (0.2)	0 (0)	0.965 (0.94)
Sim_ref	Similarity with all image references in a contest. This represents averaged value for pairwise structural similarities of each submission in a contest with each image reference provided by a contest holder	0.417 (0.41)	0.227 (0.225)	0 (0)	1 (0.97)
Sim_high_star	Similarity with a prior most recent highly-rated (4 or 5 star) submission made by another designer	0.557 (0.5)	0.276 (0.237)	0 (0)	1 (1)
Sim_synthes	Similarity with a synthesis of 2 prior highly-rated images	0.62 (0.55)	0.262 (0.247)	0 (0)	0.979 (0.98)
Sub_order	Order of each submission in each contest. Since this variable is right-skewed, I use a log transformation	107.6 (105)	129.53 (124)	1 (1)	2908 (1322)
First_sub_order	Order of first submission by each designer in each contest. Since this variable is right-skewed, I use a log transformation	78.08 (74)	117.47 (108.8)	1 (1)	2769 (1332)
Experience	Experience of each designer on the platform calculated as the number of days since joining the platform. Since this variable is right-skewed, I use a log transformation	1117 (1449)	978 (1067)	0 (4)	5836 (5637)
Contests_won	Number of prior contests won by a designer. Since	23.6 (34.7)	43.56 (60.12)	0 (0)	547 (547)

	this variable is right-skewed, I use a log transformation. This variable is highly correlated (0.63) with “experience” variable, thus I use experience variable in the main models and confirm the results with the contests_won variable in the robustness checks				
Star	A star rating for each submission. Note: the rating was provided for 186,335 submissions in open contests and for 30,975 submissions in “blind” contests	2.697 (2.61)	1.293 (1.27)	1 (1)	5 (5)

Note. Data are shown for open contests outside parentheses and for “blind” contests inside parentheses.

Table 3.2. Effect of structural similarities on the probability of winning in models with contest fixed effects

Model Independent var.	Model 1	Model 2	Model 3	Model 4
	Open contests	“Blind” contests	Open contests combined similarities	“Blind” contests combined similarities
Sub_order_log	1.82*** (0.0412)	2.06*** (0.194)	2.923*** (0.145)	3.65*** (0.513)
Star	1.9*** (0.0273)	2.12*** (0.132)	1.894*** (0.068)	2.29*** (0.238)
Experience_log	0.155*** (0.0193)	0.215** (0.085)	0.153*** (0.046)	0.11 ^{ns} (0.135)
Sim_insp	2.35*** (0.42)	1.62 ^{ns} (1.08)	3.37*** (1.27)	3.072 ^{ns} (3.67)
Sim_insp_squared	-3.265*** (0.457)	-2.498** (1.185)	-3.899*** (1.37)	-3.76 ^{ns} (4.031)
Sim_ref	1.118* (0.68)	2.3 ^{ns} (1.76)	-1.024 ^{ns} (1.19)	-1.28 ^{ns} (3.23)
Sim_ref_squared	-1.933*** (0.8)	-4.036** (1.99)	0.365 ^{ns} (1.33)	0.85 ^{ns} (3.5)
Sim_high_star	1.57*** (0.387)	1.85 ^{ns} (1.674)	1.291* (0.777)	0.637 ^{ns} (2.17)
Sim_high_star_squared	-1.84*** (0.374)	-2.39 ^{ns} (1.605)	-1.34* (0.767)	-1.21 ^{ns} (2.21)
Sim_synthes	8.525** (3.834)	3.069 ^{ns} (6.23)	7.846* (4.41)	11.67 ^{ns} (9.372)
Sim_synthes_squared	-8.31** (3.48)	-3.5 ^{ns} (5.45)	-7.62** (3.81)	-3.75 ^{ns} (2.51)
Sample size	415,258 for inspirations, 144,288 for references, and 242,637 for prior highly-rated most recent submissions, 987 for synthesized images	70,535 for inspirations, 24,656 for references, 10,621 for prior highly-rated most recent submissions, 291 for synthesized images	74,290 for all combined except synthesized, 857 for synthesized images	9,908 for all combined except synthesized, 262 for synthesized images

Note. *** - p-value <0.01; ** - p-value<0.05; * - p-value<0.1; ns – not significant.

Table 3.3. Timing of entry by experience in models with contest fixed effects

Dependent var. Independent var.	First_sub_order_log (in open contests)	First_sub_order_log (in “blind” contests)
Experience_log	0.0518*** (0.00279)	0.05845*** (0.007225)
Constant	3.28*** (0.0186)	3.162*** (0.05)
Sample size	148,854	25,266

Note. *** - p-value <0.01; ** - p-value<0.05; *- p-value<0.1; ns – not significant.

Table 3.4. Effects of experience in open and “blind” contests in models with contest fixed effects

Dependent var. Independent var.	Experience_log (in open contests)	Contests_won_log (in open contests)	Experience_log (in “blind” contests)	Contests_won_log (in “blind” contests)
Smilarity_ref_log	0.03 ^{ns} (0.622)	0.103 ^{ns} (0.955)	0.178 ^{ns} (0.4)	-0.179 ^{ns} (1.5)
Sim_insp_log	-1.1** (0.535)	-0.39*** (0.143)	-1.47** (0.711)	-0.687** (0.34)
Sim_high_star_log	-0.817*** (0.2077)	-0.838*** (0.238)	-0.946 ^{ns} (0.735)	-0.167 ^{ns} (0.589)

Note. *** - p-value <0.01; ** - p-value<0.05; *- p-value<0.1; ns – not significant. Logs are used for interpretability.

Table 4.1. Descriptive statistics for the main variables in the dataset

Name	Description	Mean	St. Dev.	Min	Max
Award	Contest award amount in US dollars	156.93	137.59	38	2,284
Winner	A winner variable dummy indicating whether a submission by a designer is a winning submission in a contest				
After	A dummy variable indicating whether a date of a contest belongs to the period after AI introduction				
Sub_order	Order of each submission in each contest	68.54	64.27	1	494
Experience	Experience of each designer on the platform calculated as the number of days since joining the platform	323.7	269.9	0	1964
Star	A star rating for each submission	2.87	1.26	1	5

Table 4.2. Confusion matrix for 5 emotions

Predicted Emotion Real Emotion	Amuse ment	Awe	Content ment	Excite ment	Sadness	Row Accuracy (%)
Amusement	65	14	8	45	7	46.7
Awe	2	185	4	2	3	94.4
Contentment	2	15	148	12	21	74.7
Excitement	29	10	17	143	19	65.6
Sadness	3	5	41	12	161	72.5

Note. The results are shown for the 20% test set (973 images out of 4,865)

Table 4.3. Descriptive statistics for the variables related to emotions and complexity

Name	Description	Mean	St. Dev.	Min	Max
Complexity	A spatial information complexity measure for each design image	0.887	0.741	0.1	2.975
Emotions_binary	A dummy variable indicating whether a design image has emotional content				
Amusement	One of the five emotions predicted in a design image by a deep learning model	1.35	1.65	0.01	53.7
Awe	One of the five emotions predicted in a design image by a deep learning model	0.319	0.99	0.003	75.69
Contentment	One of the five emotions predicted in a design image by a deep learning model	1.75	1.91	0.01	55.7
Excitement	One of the five emotions predicted in a design image by a deep learning model	3.63	3.38	0.04	59.2
Sadness	One of the five emotions predicted in a design image by a deep learning model	3.6	3.71	0.009	69.2

Table 4.4. Comparison of lower-tier and higher-tier logo contest task descriptions before and after the launch of AI

Type of a contest	Abstract req. before AI	Abstract req. after AI	Specific req. before AI	Specific req. after AI	Brand emotions (sentiment) before AI	Brand emotions (sentiment) after AI
Lower-tier logo contests	6.5	7.07	7.85	7.67	0.659	0.692
Higher-tier logo contests	8.9	9.02	9.39	9.5	0.722	0.725

Table 4.5. Propensity score matching model for designers' switching behaviors after the AI launch

Dep. variables are binary variables for switchers	Cross-tier designers versus focused designers	Cross-category designers versus focused designers
Proportion of higher-tier logo contests before AI	0.059*** (0.0044)	
Proportion of non-logo contests before AI		0.209*** (0.033)
Sample size	181,182	185,068

Note. *** - p-value<0.01; ** - p-value<0.05; * - p-value<0.1; ns – not significant

Table 4.6. Comparison of the number of contests that each designer participates in per day before and after AI for 3 groups of successful designers as compared to the unsuccessful designers

Dep. variable is the number of contests	Focused designers	Cross-tier designers	Cross-category designers	Focused designers (PSM)	Cross-tier designers (PSM)	Cross-category designers (PSM)
Unsuccessful designers after AI	0.551*** (0.0162)	0.524*** (0.0213)	0.563*** (0.06)	0.5482*** (0.016)	0.472*** (0.0215)	0.535*** (0.047)
Successful designers before AI	-0.179*** (0.0365)	-0.3086*** (0.0477)	-0.126 ^{ns} (0.1)	-0.176*** (0.0367)	-0.314*** (0.048)	-0.274*** (0.057)
Successful designers after AI	-0.123*** (0.0469)	-0.312*** (0.059)	-0.0812 ^{ns} (0.097)	-0.115*** (0.0471)	-0.312*** (0.057)	-0.187*** (0.0625)
Constant	4.1731*** (0.0092)	4.1*** (0.0129)	4.61*** (0.0159)	4.173*** (0.0091)	4.11*** (0.0128)	4.62*** (0.016)
Designer fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Sample size	73,083	73,091	45,522	62,127	61,933	39,937

Note. *** - p-value<0.01; ** - p-value<0.05; * - p-value<0.1; ns – not significant. Omitted group is “unsuccessful designers before AI”

Table 4.7. Comparison of the number of re-submissions per user per contest before and after AI for 3 groups of successful designers as compared to the unsuccessful designers

Dep. variable is the number of re-submissions	Focused designers	Cross-tier designers	Cross-category designers	Focused designers (PSM)	Cross-tier designers (PSM)	Cross-category designers (PSM)
Unsuccessful designers after AI	-1.539*** (0.0235)	0.054 ^{ns} (0.05)	0.097 ^{ns} (0.182)	-1.526*** (0.0237)	0.182*** (0.051)	0.185 ^{ns} (0.184)
Successful designers before AI	-0.543*** (0.056)	4.014*** (0.104)	-0.594 ^{ns} (1.38)	-0.541*** (0.057)	4.117*** (0.107)	-0.584 ^{ns} (1.399)
Successful designers after AI	3.12*** (0.0663)	5.755*** (0.116)	2.733** (1.38)	3.15*** (0.0668)	5.738*** (0.118)	2.79** (1.397)
Constant	5.33*** (0.014)	3.577*** (0.0348)	4.298*** (0.963)	5.334*** (0.0145)	3.576*** (0.0351)	4.282*** (0.97)
Designer fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Sample size	73,083	73,091	45,522	61,134	60,827	38,717

Note. *** - p-value<0.01; ** - p-value<0.05; * - p-value<0.1; ns – not significant. Omitted group is “unsuccessful designers before AI”

Table 4.8. Comparison of presence of emotional content before and after the AI launch for 3 groups of successful designers

Dep. variable is emotions binary	Focused successful designers	Cross-tier successful designers	Cross-category successful designers	Focused successful designers (PSM)	Cross-tier successful designers (PSM)	Cross-category successful designers (PSM)
After	0.256*** (0.029)	0.370*** (0.038)	0.577*** (0.035)	0.263*** (0.042)	0.353*** (0.06)	0.557*** (0.049)
Constant	1.205*** (0.113)	1.483*** (0.117)	1.37*** (0.14)	1.19*** (0.114)	1.469*** (0.116)	1.358*** (0.138)
Designer fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Sample size	29,072	21,773	20,953	26,659	19,639	18,627

Note. *** - p-value<0.01; ** - p-value<0.05; * - p-value<0.1; ns – not significant

Table 4.9. Comparison of complexity before and after the AI launch for 3 groups of successful designers

Dep. variable is complexity	Focused successful designers	Cross-tier successful designers	Cross-category successful designers	Focused successful designers (PSM)	Cross-tier successful designers (PSM)	Cross-category successful designers (PSM)
After	0.058*** (0.0095)	0.107*** (0.0082)	0.137*** (0.01)	0.0599*** (0.0117)	0.109*** (0.0172)	0.131*** (0.014)
Constant	0.747*** (0.0564)	0.977*** (0.056)	0.837*** (0.077)	0.7472*** (0.0565)	0.975*** (0.057)	0.838*** (0.076)
Designer fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Sample size	29,072	21,773	20,953	26,659	19,639	18,627

Note. *** - p-value<0.01; ** - p-value<0.05; * - p-value<0.1; ns – not significant

Table 4.10. Comparison of five emotions before and after the AI launch for 3 groups of successful designers

Dep. variable is After AI	Focused successful designers	Cross-tier successful designers	Cross-category successful designers	Focused successful designers (PSM)	Cross-tier successful designers (PSM)	Cross-category successful designers (PSM)
Amusement	-0.0421 ^{ns} (0.035)	0.022 ^{ns} (0.024)	-0.011 ^{ns} (0.022)	-0.029 ^{ns} (0.025)	0.0229 ^{ns} (0.0244)	-0.03 ^{ns} (0.021)
Awe	0.009 ^{ns} (0.034)	-0.011 ^{ns} (0.047)	-0.03 ^{ns} (0.053)	-0.016 ^{ns} (0.035)	-0.0105 ^{ns} (0.046)	-0.038 ^{ns} (0.045)
Contentment	-0.012 ^{ns} (0.0136)	-0.0279 ^{ns} (0.02)	0.0178 ^{ns} (0.0193)	-0.0265 ^{ns} (0.0135)	-0.0286 ^{ns} (0.021)	-0.0183 ^{ns} (0.0169)
Excitement	0.177 ^{***} (0.00739)	0.331 ^{***} (0.043)	0.517 ^{***} (0.00804)	0.17 ^{***} (0.00671)	0.337 ^{***} (0.033)	0.584 ^{***} (0.0094)
Sadness	-0.0002 ^{ns} (0.00747)	-0.0194 ^{ns} (0.014)	-0.0099 ^{ns} (0.011)	-0.007 ^{ns} (0.0072)	-0.0192 ^{ns} (0.019)	-0.0153 ^{ns} (0.009)
Constant	-1.658 ^{***} (0.239)	-1.49 ^{***} (0.273)	-0.626 ^{***} (0.277)	-1.641 ^{***} (0.238)	-1.48 ^{***} (0.272)	-0.645 ^{***} (0.278)
Designer fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Sample size	29,072	21,773	20,953	26,659	19,639	18,627

Note. *** - p-value<0.01; ** - p-value<0.05; * - p-value<0.1; ns – not significant

Table 4.11. Comparison of presence of emotional content before and after the AI launch for 3 groups of unsuccessful designers

Dep. variable is emotions binary	Focused unsuccess. designers	Cross-tier unsuccess. designers	Cross-category unsuccess. designers	Focused unsuccess. designers (PSM)	Cross-tier unsuccess. designers (PSM)	Cross-category unsuccess. designers (PSM)
After	0.034 ^{ns} (0.023)	0.0143 ^{ns} (0.0297)	0.054 ^{ns} (0.039)	0.052 ^{ns} (0.037)	-0.024 ^{ns} (0.047)	0.0722 ^{ns} (0.059)
Constant	1.35 ^{***} (0.088)	1.55 ^{***} (0.114)	1.49 ^{***} (0.128)	1.347 ^{***} (0.0885)	1.54 ^{***} (0.1142)	1.495 ^{***} (0.129)
Designer fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Sample size	60,044	31,687	18,934	53,859	29,056	16,605

Note. *** - p-value<0.01; ** - p-value<0.05; * - p-value<0.1; ns – not significant

Table 4.12. Comparison of complexity before and after the AI launch for 3 groups of unsuccessful designers

Dep. variable is complexity	Focused unsucc. designers	Cross-tier unsucc. designers	Cross-category unsucc. designers	Focused unsucc. designers (PSM)	Cross-tier unsucc. designers (PSM)	Cross-category unsucc. designers (PSM)
After	-0.0059 ^{ns} (0.0057)	-0.0173 ^{ns} (0.0121)	-0.0084 ^{ns} (0.0059)	-0.012 ^{ns} (0.008)	-0.0095 ^{ns} (0.013)	-0.0057 ^{ns} (0.0123)
Constant	0.892 ^{***} (0.049)	0.951 ^{***} (0.061)	0.972 ^{***} (0.0123)	0.8923 ^{***} (0.0491)	0.952 ^{***} (0.0611)	0.973 ^{***} (0.0123)
Designer fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Sample size	60,044	31,687	18,934	53,859	29,056	16,605

Note. *** - p-value<0.01; ** - p-value<0.05; * - p-value<0.1; ns – not significant


Table 4.13. Comparison of five emotions before and after the AI launch for 3 groups of unsuccessful designers

Dep. variable is After AI	Focused unsucc. designers	Cross-tier unsucc. designers	Cross-category unsucc. designers	Focused unsucc. designers (PSM)	Cross-tier unsucc. designers (PSM)	Cross-category unsucc. designers (PSM)
Amusement	-0.0164 ^{ns} (0.017)	-0.0038 ^{ns} (0.0056)	0.0054 ^{ns} (0.0074)	-0.0173 ^{ns} (0.015)	-0.023 ^{ns} (0.017)	-0.0119 ^{ns} (0.0128)
Awe	0.016 ^{ns} (0.0103)	-0.0004 ^{ns} (0.0092)	0.0093 ^{ns} (0.0094)	0.0112 ^{ns} (0.0155)	0.005 ^{ns} (0.03)	-0.014 ^{ns} (0.023)
Contentment	-0.00096 ^{ns} (0.0055)	-0.0011 ^{ns} (0.0047)	-0.0044 ^{ns} (0.0059)	-0.01 ^{ns} (0.0079)	-0.0065 ^{ns} (0.012)	-0.008 ^{ns} (0.01)
Excitement	0.0016 ^{ns} (0.0032)	-0.00032 ^{ns} (0.0027)	0.005 ^{ns} (0.0037)	-0.0079 ^{ns} (0.0061)	-0.0052 ^{ns} (0.0069)	0.0078 ^{ns} (0.00593)
Sadness	-0.00444 ^{ns} (0.0031)	0.00142 ^{ns} (0.0024)	-0.00086 ^{ns} (0.003)	-0.0101 ^{ns} (0.0072)	0.006 ^{ns} (0.00593)	-0.0068 ^{ns} (0.0054)
Constant	-1.2 ^{***} (0.132)	-0.889 ^{***} (0.137)	-0.937 ^{***} (0.124)	-1.22 ^{***} (0.133)	-0.912 ^{***} (0.141)	-0.937 ^{***} (0.124)
Designer fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Sample size	60,044	31,687	18,934	53,859	29,056	16,605

Note. *** - p-value<0.01; ** - p-value<0.05; * - p-value<0.1; ns – not significant

Figures

Figure 2.1. Content similarity with top competitors



Note. Numbers represent pairwise cosine similarity of content.

Figure 2.2. Examples of top competitors with one outlier (Panel A) and two outliers (Panel B)




Figure 2.3. Hierarchy of firm-tweet categories

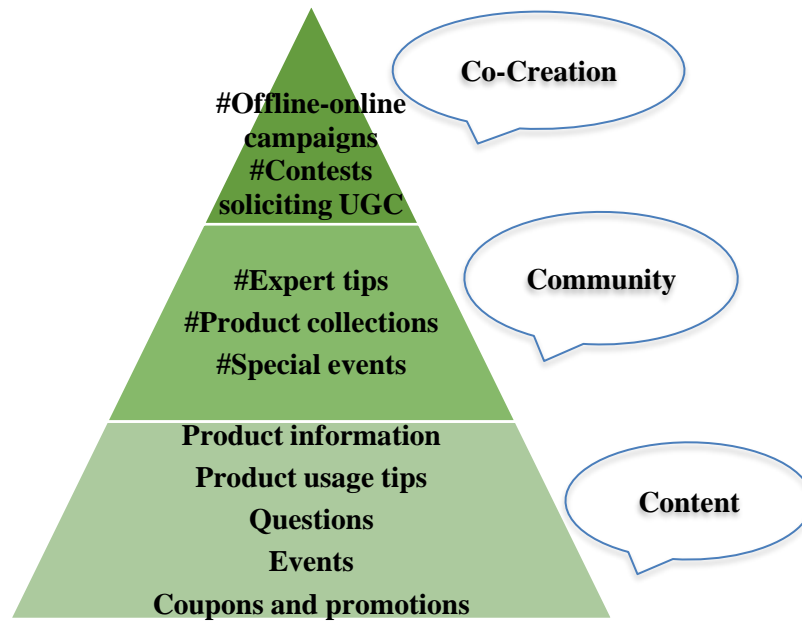


Figure 3.1. Structural similarity index example of similar images (similarity is 0.92)



A



B

Figure 3.2. Structural similarity index example of dissimilar images (similarity is 0.14)



C



D

Figure 3.3. Two original highly-rated images (A and B) and their synthesized images (C and D)




Figure 3.4. Source images A and B, synthesized images and corresponding winning images



Figure 4.1. Examples of a simple logo (Image A, complexity score: 0.235) and a complex logo (Image B, complexity score: 2.05)



A



B

Figure 4.2. Examples of logos with emotional content (image C) and without emotional content (image D)

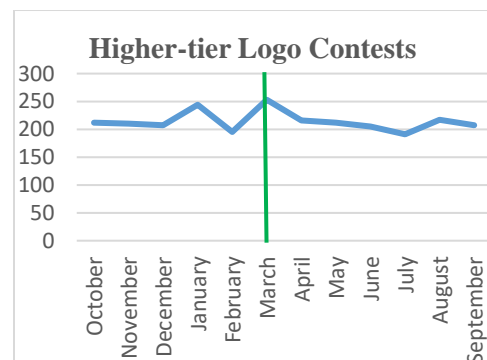



C



D

Figure 4.3. The number of lower-tier and higher-tier logo contests before and after AI



Appendices

Table A1. Examples of cosine similarity scores for tweets within content category and across categories

Tweet 1	Tweet 2	Within category	Cosine similarity
The Me on GNC #contest ends in just 3 days! Enter your photo/video before July 1st at http://t.co/xBcnCLMikc & you could #win \$25,000!	Share your Best #DadAdvice to win our Dad's Day Contest! Win 1 of 6 prizes! Official contest: http://t.co/N5EGtSL3vJ	yes	0.238
CAPTION CONTEST! How fresh is your neck? Submit your entry in our caption contest to receive your free #DXLTIE -> http://t.co/8tBzVBorU0...	Share your Best #DadAdvice to win our Dad's Day Contest! Win 1 of 6 prizes! Official contest: http://t.co/N5EGtSL3vJ	yes	0.192
Share your #LoveGUESS ?? Snap a pic w/ #LoveGUESS #GUESSContest + @GUESS for a chance to win https://t.co/8zSnW9qKPP https://t.co/aWswOHQW9t	Share your Best #DadAdvice to win our Dad's Day Contest! Win 1 of 6 prizes! Official contest: http://t.co/N5EGtSL3vJ	yes	0.15
The Me on GNC #contest ends in just 3 days! Enter your photo/video before July 1st at http://t.co/xBcnCLMikc & you could #win \$25,000!	What is a #LVMHday? Take a video tour during the first event organized end of 2016 at @HECParis. #LVMHtalents... https://t.co/4ZvrCAaGGx	no	0.073
Share your Best #DadAdvice to win our Dad's Day Contest! Win 1 of 6 prizes! Official contest: http://t.co/N5EGtSL3vJ	Imitation is the best form of flattery! #LikeMotherLikeDaughter #SKECHERSstyle #SKECHERSDemiStyle http://t.co/gI8EheFBnb	no	0.072
Share your Best #DadAdvice to win our Dad's Day Contest! Win 1 of 6 prizes! Official contest: http://t.co/N5EGtSL3vJ	Don't let your greens go, blend them up as a tasty start to the day. @plantstrongveg has 3 must try recipes!... https://t.co/ZTZDV7av42	no	0.042

Table A2. Description of tweet categories and examples of tweets in each category

Category	Description	Sample of tweets
Product information	This category of tweets is used by firms to introduce products and provide some information to online users. For example, Nike announced that “The Air Jordan 10 Retro “Double Nickel” is released online. Another more general example is to showcase a new product and to encourage online users to buy it, or to highlight that some popular items have limited stocks.	<p>Take your sweet potatoes to the next level with these Baked Sweet Potato Fries. http://t.co/sq2knEYDR2</p> <p>Our essentials tees + palazzo pants + chambray + wide brim hat = YES. http://t.co/HCzddBzz6R http://t.co/fvJirpSm2W</p> <p>We have a couple new colorways of the Nike Roshe Run available here: http://t.co/JgHXODyB6g @TeamRoshe http://t.co/xOefiuDRzX</p> <p>New season, new beauty products! Our CFT Camomile Sumptuous Cleansing Butter soothes your skin> http://t.co/gq3Q1AzpC9 @examinercom</p> <p>We're swooning over the McKinley Lace Dress! http://t.co/ROfvS7ZqZM http://t.co/k29KXatdXB</p>
Product usage tips	This category of tweets is related to specific tips on how to use a firm’s products. For example, firms share tips on how to create a gallery wall with one of its products, or tips related to home furnishings with a firm’s products, or tips related to how to create the best gift that will include a firm’s product.	<p>Don't let your greens go, blend them up as a tasty start to the day. @plantstrongveg has 3 must try recipes!... https://t.co/ZTZDV7av42</p> <p>Dinning room lighting a little outdated? Spruce it up with this easy how-to tutorial. http://t.co/GxXr363Imh</p> <p>Make the most of your outdoor space - tips and tricks to make outdoor living a breeze. https://t.co/jZtHnkq4LV https://t.co/pZuEXb7dth</p> <p>What a great idea...everything you need to personalize your own DIY area rug!!!! Can't wait to try! :-) http://t.co/1HmsPwYGmP</p> <p>Read our how to wear guide for our monk strap brogues "Busby Jazz": http://t.co/5pUHN557YB http://t.co/IQVzMwGjNZ</p>
Questions	This category of tweets asks online users open-ended	What summer accessory are you most excited to wear?

	<p>questions with a blank or without blanks. For instance, firms could tweet “Finish the saying: If the shoe fits, _____”, or “What are your weekend plans?”, or “What is the one crafting supply you couldn't live without?”.</p>	<p>Fill in the blank: My kids favorite movie_____</p> <p>What's your favorite way to wear fragrance during the summer?</p> <p>If you could get a birthday shout-out from someone famous, who would it be?</p> <p>Anyone hit the slopes this weekend?</p>
Events	<p>This category of tweets is related to events, where firms share information about upcoming events. For example, firms tweet about new store opening events, Q&A events related to panel discussions with experts, or events related to some live shows.</p>	<p>If you're in the Dallas area, you can meet @JoeyGraceffa on Saturday! See more details at https://t.co/oOFkzzk5fg http://t.co/kNAJR0J0AD</p> <p>Join us for Part 2 of our "Healthy in a Hurry" Series this weekend! Times vary- Call your local store for RESV details: http://t.co/8W1rPVxi</p> <p>Heading out to the SFentrepreneur.com meet up at Action Theatre. Hope to see you there.</p> <p>Visit us at the Southern Ideal Home Show in Raleigh starting today. http://t.co/Ceg7BddZ</p> <p>Have you signed up for the Build & Grow workshop this weekend? We hear it's gong to be spooktacular! http://t.co/lJ0fydYo</p>
Coupons and promotions	<p>This category of tweets is used to share information about online and offline coupons and promotions available in all stores, or in specific offline locations or exclusively online. For example, a firm might use the following tweet: “Save 20% on all dining tables and chairs, buffets, china cabinets, & more! Limited time only!”. Another example is “Saturday steal: spend \$125 and get 40% off your highest priced item.”. Additionally, a tweet might contain an external link with more information.</p>	<p>Wednesday Special: Get \$10 off when you spend \$75 online. Today through 4pm CT! http://t.co/17pNkPTn0p</p> <p>Save 25% on all Ceramic Bakeware and Dutch Ovens. https://t.co/EwsckMCB5y https://t.co/qXwUA4ABPP</p> <p>TODAY ONLY – Save 40% off Nutella! Sale is valid 4/24/13 only, limit 4 per customer while supplies lasts. http://t.co/jvDPbMO6QQ</p> <p>Sweaters, Shirts, Tees! Buy One, Get One 50% Off! Shop Tops > http://t.co/HoB88NEgB1</p>

		<p>Saturday steal: spend \$125 and get 40% off your highest priced item. https://t.co/RQUrXmBbh3 https://t.co/MG37iQ14D0</p>
#Expert tips	<p>This category of tweets is related to online collaboration with influential social media users. These influential users provide tips related to style/look/products. For example, firms collaborate with popular bloggers or social media stars in an attempt to engage their followers and invite those followers to participate in the discussions.</p>	<p>Pull your #hair back with a pretty bow for an added touch of special. #AskHairGenius for more style tips. http://t.co/2S4PEmtwSL</p> <p>7 Ways To Refresh Your Home Office in our Tips & Ideas! https://t.co/7XgJAQ9Ap7 #homeoffice #pbstyletip #design https://t.co/s9YI7unUpD</p> <p>Tip of the Day: The tuck in is key... #DXLTips #GuyStyle #SweaterWeatherRules https://t.co/1yZNZtX9bO</p> <p>Laidback luxe has reached new heights—the track pant is a street style star. #StyleTip</p> <p>Tip: To wear your @LoxStudio extensions in an up-do, flip your hair over & snap them in the opposite direction, then style. #BeautyReport</p>
#Product collections	<p>This category of content promotes specific product collections under a hashtag that online users could track. For example, Puma introduced a new collection using the #FENTYxPUMA hashtag (rap-up.com 2016).</p>	<p>Comfort, quality and innovation: last day at @imm_cologne for the new collection of #NatuzziEditions. Discover it... https://t.co/PZ0HSsvwdf</p> <p>Celebrating the new #CoachAndRodarte collection with #ChloeGraceMoretz last night in L.A. https://t.co/JMxZqdgw7J</p> <p>Last Chance to shop the current #CRAFTBYWORLDMARKET collection! New collection tomorrow. http://t.co/aNwFqooUum http://t.co/S6tow4wEFg</p> <p>Preview of the SS13 and #ProGreenxPUMA collections. @professorgreen http://t.co/dbje56hT</p> <p>Lucky You...Lucky Brand 's Fall collection has arrived, check it out now! #DXL #luckybrand</p>

		https://t.co/QuzhOdDsIv https://t.co/s4Gexw8ybv
#Special events	<p>This category of tweets includes sponsorship for a series of events. An example is the #nmmakesomenoise (blog.depict.com 2015) campaign by Neiman Marcus, which is related to a series of events that “feature thought-provoking panel discussions with female noisemakers in technology, fashion, film, music, food, art, and business”.</p>	<p>It's that time of year! Join us today for our #GUESSHoliday Events in stores! Info at https://t.co/4VyaIR4s7T https://t.co/ca0OVEOR9l</p> <p>Around The Store: Talking trends with @harpersbazaarus' @Avril_at_BAZAAR at NM Westchester. #NMevents http://t.co/CGfB5mUcdo</p> <p>Today! We're having #DVFlovesGap events in #NYC #LA #SF & more. http://t.co/BuBN60KsBN Bring your #GapKids for special activities. Cc @DVF</p> <p>Join us for #FNO tomorrow! To find an event in your area, click here: http://t.co/ct306dQj</p> <p>What is a #LVMHday? Take a video tour during the first event organized end of 2016 at @HECParis. #LVMHtalents... https://t.co/4ZvrCAaGGx</p>
#Offline-online campaigns	<p>This category of tweets is related to firms' cross-channel marketing efforts. Firms with offline marketing campaigns often involving celebrities, use Twitter to not only solicit creative ideas from users but also invite users to follow celebrities' example and contribute content. Typically, these campaigns originate offline, and are promoted in news sources as campaigns with a specific Twitter hashtag.²⁶ Thus, firms invest resources to closely integrate their marketing and promotions across offline as well as social media platforms.²⁷</p>	<p>Be anything. Do everything. Love anyone. Play nice. Make peace. Stay true. #WeAllCan @YaraShahidi @guggle23... https://t.co/EOTyyM0XGm</p> <p>Show us a little love and we'll show you a little luck! #GearUpForGreat #HappyStPatricksDay https://t.co/hWnihXvuub</p> <p>Imitation is the best form of flattery! #LikeMotherLikeDaughter #SKECHERSstyle #SKECHERSDemiStyle http://t.co/gI8EheFBnb</p>

²⁶ I search offline press reports to confirm that the tweets in this category involve firms' offline campaigns and that these offline campaigns include the Twitter hashtags as well.

²⁷ Twitter picked Nike's campaign starring Colin Kaepernick as the most creative brand campaign in 2018 (campaignlive.co.uk 2019). That offline-online campaign featured quarterback Colin Kaepernick who opposed the tradition of standing to the national anthem as a protest against racial injustice. The

		<p>You are the sum of all your training. #RuleYourself #IWILL https://t.co/xwupw6U5kM</p> <p>Run straight into the holidays. Keep it merry and bright every time you lace up. #ForeverFaster http://t.co/5YBxqHpYMn</p>
#Contests soliciting UGC	<p>This category consists of tweets relating to online contests asking online users to upload user-generated content in the form of advice, design suggestions, photos, photo captions, short stories (sproutsocial.com 2016) or videos involving a firm's products. As part of these contests, users are also encouraged to vote for other users' content, thus adding additional interactivity to the campaign. The contests are not only funny and engaging, but also allow harnessing innovative inputs from online users and using those inputs for content curation.</p>	<p>Enter our #PinterestContest through 6/18 for a chance to win a \$5k room makeover! http://t.co/JmkoK9asnB http://t.co/WAU2hrftMR</p> <p>Share your #LoveGUESS ?? Snap a pic w/ #LoveGUESS #GUESSContest + @GUESS for a chance to win https://t.co/8zSnW9qKPP https://t.co/aWswOHQW9t</p> <p>The Me on GNC #contest ends in just 3 days! Enter your photo/video before July 1st at http://t.co/xBcnCLMikc & you could #win \$25,000!</p> <p>CAPTION CONTEST! How fresh is your neck? Submit your entry in our caption contest to receive your free #DXLTIE -> http://t.co/8tBzVBorU0...</p> <p>Share your Best #DadAdvice to win our Dad's Day Contest! Win 1 of 6 prizes! Official contest: http://t.co/N5EGtSL3vJ</p>

tweet "Believe in something, even if it means sacrificing everything. #JustDoIt" sparked great conversation and received millions of "likes". Nike integrated that campaign with a hashtag #JustDoIt that uses a slogan "Just do it" with a long history in the Nike's brand.

Section A1. Methodology to Calculate the Approximate “Date of Following” a Focal Firm by Each Follower

Following the methodology proposed by Bruns et al. (2014), to calculate the approximate date when a follower starts following a focal firm, for any user u , one can calculate the earliest possible date at which they could have followed the focal firm, as the most recent account creation date encountered for accounts 1 through u :

$$\text{followdate}_{\min}(u) = \max(\text{creationdate}(1): \text{creationdate}(u)) \quad (1)$$

where 1 – is the earliest follower. In essence, the algorithm looks at the earliest follower’s date of account creation, then looks at the next follower (higher in the order, i.e., started following a focal firm later than the earliest follower) and his/her date of account creation. If the date of account creation of the first follower is higher (i.e., later, for example, May 2013) than the date of account creation of the second later follower (for example, May 2012), it means that the second follower’s minimum date when he/she could have started following a focal firm is at least the same as the date of the first follower account creation (May 2013).

The described approach relies on so-called “anchors”, i.e., users who created account, for example in each quarter of the study period 2012-August 2017, and then started following a focal firm. I find that, with many followers (hundreds of thousands), the number of anchors is quite large (several thousand). Thus, this approach for calculating followers “date of following” should be accurate enough. Additionally, the mathematical proof of that method is provided in the paper by Meeder et al. (2011).

Additionally, I validate this method by comparing the estimated number of followers with the actual number of followers for a subset of 30 firms for which I can observe the changes in the actual number of followers by quarter on the website of

historical snapshots of Web pages web.archive.org. The differences between estimated number of followers and the actual number of followers for select firms are within plus/minus 10% and, thus, these differences do not affect the results when I add/subtract 10% from the estimated number of followers as a robustness check.


Section A2. Robustness Check - Comparison of Social Media and 10-K Reports

I analyze the business descriptions of the firms' 10-K reports as a robustness check of traditional competition. Traditional competitors are expected to be more similar in their 10-K content than non-competing firms. I collect a subset of 10-K annual reports for 71 firms for the period 2012-2015 for the Retail Trade sector, use business description of each report to create a TFIDF vector for each firm for each year, then use those TFIDF vectors to calculate pairwise cosine similarity for each pair of firms in the dataset, and recalculate the social media pairwise cosine similarity metric (to match the 10-K metric) for each pair of firms by year (previously it was calculated by quarter).

Figure A1 shows that the two firms represented by green nodes were more distant from their competitors based on their Twitter content in 2015, while all 4 firms were very similar to each other in 2015 based on analysis of their 10-K reports.

I perform this robustness check for all subsets of top competitors and find consistent differences between social media and 10-K reports. In other words, firms that were classified as traditional competitors by Hoover's, were found to have a high degree of similarity based on a content analysis of their 10-K reports. However, a content analysis of their Twitter content highlights firms that are dissimilar to their traditional competitors.

Figure A1. Comparison of social media similarity (Panel A) with 10-K similarity (Panel B) in 2015



Additionally, I assume that 10-K reports' business description sections' similarity reflects strength of traditional competition. To test that assumption, I check whether top competitors have higher similarity than other (non-top) competitors, and that other competitors have higher similarity than non-competitors. I find that this assumption is supported in my subset of 10-K reports. When it comes to social media, top competitors' content strategies can change from year to year, so that in some years top competitors might be closer to each other than other competitors and non-competitors, while in other years top competitors are more distant from each other compared to other competitors and non-competitors. My results are robust to these changes in firms' content strategies across time.

Section A3. Deep Learning Models Description

I train two deep Recurrent Neural Network (RNN) models with Long Short-Term Memory (LSTM) using the Keras package (keras.io) of python – one model for 5 categories with hashtags and another model for 5 categories without hashtags. To

improve the models' accuracy, I use external embeddings with the dimension 200 from GloVe (Global Vectors for Word Representation, nlp.stanford.edu), pre-trained specifically for Twitter on 2 billion tweets. The last layer uses sigmoid activation function to output a distribution of probability where each represents an input tweet being classified as the corresponding category. Further, a grid search technique is used for finding the best parameters, and the early stopping is also used to prevent overfitting. The overall first model accuracy for tweets with hashtags is 86% on the hold-out test set (20% of data, average precision is 83.7, average recall is 87.9), and 83% under the 10-fold stratified cross-validation. The overall second model accuracy for tweets without hashtags is 87% on the hold-out test set (20% of data, average precision is 84.9, average recall is 88.8), and 85% under the 10-fold stratified cross-validation. In the 10-fold cross-validation setting, the algorithm runs 10 times, and each time each training dataset uses 90% of random labelled tweets, while the test dataset includes the remaining 10% of random labelled tweets.

Since the category “#offline-online campaigns” represents tweets with hashtags related to some offline campaigns, I use additional robustness check to make sure those hashtags appear in the news. The deep learning model identified about 41,156 tweets as potentially belonging to the category “#offline-online campaigns”. I use an automated Google search (top 10 search results) to check whether each hashtag in each tweet appears in the following press-releases websites: prnewswire.com, businesswire.com, prlog.org, pr.com, and prweb.com. I search each hashtag and check whether top 10 results of the search (i.e., links) contain those press-releases websites. The top 10 results were chosen based on a manual pilot test search. If I use a less

restrictive search (all news as opposed to just press-releases), then the number of confirmed tweets in that category is slightly higher, but that number has normalized engagement similar to the tweets that are confirmed in press-releases only. Out of 41,156 candidate tweets, I confirm for 13,017 tweets (Table 2.5 in the paper) that at least one hashtag appears in the press-releases in top 10 Google search results. The remaining 28,139 tweets are assigned to the category “Misc.”.

Section A4. Mediation Analysis

The findings point to a mediation process where dissimilarity affects online engagement and new followers’ acquisition rate through the usage of higher-level affordances. To test for full or partial mediation of the effect of dissimilarity on online engagement and new followers’ acquisition rate, I use the structural equation modeling method (Stata “SEM” package). I combine all 5 higher-level categories into a new variable by summing up proportions of each of the individual higher-level categories (tiers 1 and 2 in Table 2.5 of the paper) for each firm for each quarter (averaging proportions for all 5 higher-level categories would give the same result). If higher-level categories fully mediate the effect of dissimilarity on online engagement and new followers’ acquisition rate, then the direct effect of dissimilarity on those dependent variables should become non-significant when “higher-level categories” variable is included in the model as a mediator.

The conceptual mediation equation is shown below:²⁸

$$\text{sem (MV <- IV CV1 CV2) (DV <- MV IV CV1 CV2)} \quad (2)$$

²⁸ Mediation example for SEM package of Stata: <https://stats.idre.ucla.edu/stata/faq/how-can-i-do-mediation-analysis-with-the-sem-command/>

where MV refers to the mediator variable (higher-level categories); DV refers to the dependent variables (favorites, retweets or new followers); IV refers to the independent variable (similarity); CVs are covariates (the number of tweets and the number of followers). Table A3 in this Appendix illustrates the results.

Table A3 shows that higher-level categories of affordances fully mediate the effect of dissimilarity on online engagement and new followers' acquisition rate. The direct effect of similarity becomes non-significant in the presence of the mediator (i.e., "higher-level categories", which has a statistically significant direct effect) in the model. The direct effect of similarity on the use of higher-level categories (not reported in the Table A3) is negative "–100.7" with the *p*-value of less than 0.001, which means that dissimilarity is associated with the usage of higher-level categories. As a robustness check, I operationalize higher-level categories as counts of tweets in specific categories for each firm for each quarter (not as proportions), and the results are consistent.

Table A3. Structural Equation Modeling mediation results

Depend. Var.	Direct effect of similarity	Indirect effect of similarity	Total effect of similarity	Direct effect of higher-level categories	Indirect effect of higher-level categories	Total effect of higher-level categories
Favorites	-13,455 (<i>p</i> = 0.126)	-5,802 (<i>p</i> =0.011)	-19,26 (<i>p</i> = 0.03)	57.61 (<i>p</i> = 0.001)	No Path	57.61 (<i>p</i> = 0.001)
Retweets	-6,325.5 (<i>p</i> = 0.07)	-1,610.7 (<i>p</i> =0.021)	-7,936.3 (<i>p</i> =0.016)	15.99 (<i>p</i> = 0.021)	No Path	15.99 (<i>p</i> = 0.021)
New followers	-13,787.6 (<i>p</i> =0.078)	-12,871 (<i>p</i> =0.000)	-26,658.1 (<i>p</i> =0.000)	127.8 (<i>p</i> =0.00)	No Path	127.8 (<i>p</i> =0.000)

Notes. Estimations are performed with robust standard errors. Estimations with bootstrapping are consistent.

Table A4. Effect of structural similarities on the probability of winning in open contests

Model Independent var.	Model 1	Model 2	Model 3	Model 4	Model 5
	Only references are provided	Only inspirations are provided	Only high-star submissions are available	User and contest fixed effects	Minimal textual descript.
Sub_order_log	1.559*** (0.46)	1.659*** (0.098)	2.52*** (0.238)	3.476*** (0.231)	2.64*** (0.192)
Star	1.925*** (0.55)	1.578*** (0.0724)	2.17*** (0.0724)	1.83*** (0.1)	1.725*** (0.085)
Experience_log	0.29 ^{ns} (0.173)	0.16*** (0.044)	0.327*** (0.0964)	Omitted in user and contest FE	0.135** (0.065)
Sim_insp		2.442** (1.019)		4.68** (2.19)	2.875* (1.61)
Sim_insp_squared		-3.7*** (1.09)		-4.041* (2.29)	-4.31** (1.75)
Sim_ref	3.195 ^{ns} (4.85)			-3.81* (2.13)	-0.509 ^{ns} (1.22)
Sim_ref_squared	-4.598 ^{ns} (5.45)			3.38 ^{ns} (2.29)	-0.37 ^{ns} (1.37)
Sim_high_star			2.2* (1.31)	2.455** (1.206)	1.19* (0.71)
Sim_high_star_squared			-2.493* (1.33)	-2.878** (1.2)	-1.37* (0.77)
Sim_synthes				7.14* (4.247)	9.345* (5.41)
Sim_synthes_squared				-6.37* (3.745)	-11.129* (5.81)
Sample size	1,649	36,099	15,745	12,912 for all combined similarities except synthesized , 574 for synthesized images	4,004 for all combined similarities except synthesized , 241 for synthesized images

Note. *** - p-value <0.01; ** - p-value<0.05; * - p-value<0.1; ns – not significant. In Model 1 only references are provided (no submissions with star ratings and no inspirations). In Model 2 only inspirations are provided (no submissions with star ratings and no references). In Model 3 only high-star submissions are available (no inspirations/references). Model 5 represents a model for 935 contests with minimal textual descriptions. Models 1,2,3, and 5 use contest fixed effects. Model 4 uses user (designer) and contest fixed effects.

Table A5. Relationship between experience and information synthesis in open and “blind” contests (model with contest fixed effects)

Dependent var. \ Independent var.	Log of experience (open contests)	Log of contests won (open contests)	Log of experience (“blind” contests)	Log of contests won (“blind” contests)
Sim_synthes_log (50/50)	0.0347*** (0.0133)	0.0535*** (0.0182)	0.04 ^{ns} (0.0267)	0.0384 ^{ns} (0.033)
Sim_synthes_log (70/30)	0.0338** (0.0135)	0.0516*** (0.0186)	0.0435 ^{ns} (0.0275)	0.0296 ^{ns} (0.0343)
Sim_synthes_log (30/70)	0.0317** (0.0125)	0.0521*** (0.0171)	0.0405 ^{ns} (0.0297)	0.0303 ^{ns} (0.0361)
Sim_synthes_log (90/10)	-0.123*** (0.037)	-0.114*** (0.033)	-0.132 ^{ns} (0.347)	-0.13 ^{ns} (0.398)
Sim_synthes_log (10/90)	-0.097*** (0.0172)	-0.103*** (0.017)	-0.107 ^{ns} (0.122)	-0.151 ^{ns} (0.284)

Note. *** - p-value <0.01; ** - p-value<0.05; * - p-value<0.1; ns – not significant. Values are shown for the neural style transfer method. The synthesis proportions are shown as 50/50, 70/30 etc. in the first column.

Section A5: Comparison of AI Logo Design Emotions and Complexity with Human Logo Design Emotions and Complexity

To understand the capabilities of the AI system with respect to design emotional content and complexity, I use the AI system to generate 1,078 different logos and measure emotional content and complexity of each AI-generated logo. With respect to emotions, I predict 5 emotions in AI-generated logo images and compare those with the predicted 5 emotions in human-generated logos. Interestingly, among human logos there are more logos (higher proportion, i.e., 60.2% vs 43%) with positive emotions, fewer logos with neutral emotions, and human logos have higher amusement and excitement emotion scores (higher by 36% for excitement and by 53% for amusement in absolute scores), but a lower contentment score (lower by 41.4%) as compared to the AI-generated logos. With respect to complexity, interestingly, I find that the AI system can generate logos with complexity ranging from 0.1 to 1.14, while humans can

generate logos with complexity ranging from 0.1 to 2.975. In my dataset 54.16% of designers produce at least one logo that is more complex than the AI maximum level of 1.14, while, overall, 27.8% of all human logo submissions have complexity higher the AI maximum level of 1.14.

Section A6: Effects of Emotions and Complexity on the Probability of Winning a Contest

Given the importance of emotional content and design complexity for logo designs, I evaluate whether those variables have positive effects on the probability of winning a contest.

Table A6 shows that the effect of presence of emotional content on the probability of winning a contest is positive and significant for lower-tier logo contests, higher-tier logo contests and non-logo contests in both periods (before and after the AI launch).

Table A6. Effects of emotional content and complexity on the probability of winning a contest

Dependent variable is Winner dummy	Lower-tier logo contests before AI	Lower-tier logo contests after AI	Higher-tier logo contests before AI	Higher-tier logo contests after AI	Non-logo contests before AI	Non-logo contests after AI
Complexity	0.127*** (0.045)	0.25*** (0.054)	0.354*** (0.051)	0.307*** (0.055)	0.275** (0.134)	0.15* (0.077)
Emotions_binary	0.145* (0.087)	0.384*** (0.14)	0.27*** (0.09)	0.16* (0.09)	0.457** (0.22)	0.245* (0.15)
Sub_order	0.0017*** (0.00055)	0.0037*** (0.00058)	0.0019*** (0.0007)	0.00087 ^{ns} (0.00077)	-0.0025 ^{ns} (0.00322)	-0.00434 ^{ns} (0.0294)
Star	0.54*** (0.0147)	0.617*** (0.019)	0.47*** (0.018)	0.495*** (0.019)	0.43*** (0.045)	0.49*** (0.037)
Experience	0.0004*** (0.00012)	0.0006*** (0.00011)	0.00011 ^{ns} (0.00012)	0.00055 ^{ns} (0.00074)	-0.00039 ^{ns} (0.00034)	0.00033 ^{ns} (0.00021)
Designer fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Sample size	82,306	62,345	51,887	44,096	4,200	3,618

Note. *** - p-value<0.01; ** - p-value<0.05; * - p-value<0.1; ns – not significant

Additionally, Table A6 shows that the effect of design complexity on the probability of winning is positive in all three groups of contests in both periods.

Next, I use a more granular model for 5 emotions to see which specific emotions affect the probability of winning a contest. Since the deep learning predictive model estimates probability of each emotion for each design image, I assign an emotion label to an image by a “majority vote”, i.e., by the highest probability for one of the 5 emotions. All design images that have low probability for each of the five emotions are excluded from the analysis since they either have no emotional content or might represent some other emotion beyond the five emotions of focus. As a robustness check, I also add this indicator variable for “other/neutral” content to the model and find the results to be consistent. Table A7 reports the results.

Table A7. Effect of 5 emotions (dummy variables) on the probability of winning

Dependent var. is Winner dummy	Lower-tier logo contests before AI	Lower-tier logo contests after AI	Higher-tier logo contests before AI	Higher-tier logo contests after AI	Non-logo contests before AI	Non-logo contests after AI
Amusement	0.312* (0.18)	0.316 ^{ns} (0.211)	0.196* (0.117)	0.632*** (0.196)	0.51 ^{ns} (0.5)	0.166 ^{ns} (0.51)
Awe	0.114 ^{ns} (0.358)	-0.337 ^{ns} (0.493)	-0.195 ^{ns} (0.256)	0.32 ^{ns} (0.39)	0.2 ^{ns} (1.1)	1.01 ^{ns} (0.7)
Contentment	0.187 ^{ns} (0.146)	0.115 ^{ns} (0.172)	-0.008 ^{ns} (0.099)	0.286 ^{ns} (0.177)	0.22 ^{ns} (0.53)	0.44 ^{ns} (0.48)
Excitement	0.284*** (0.086)	0.192* (0.1)	0.0292 ^{ns} (0.069)	0.342*** (0.1)	0.84*** (0.28)	0.491* (0.265)
Sub_order	0.0019** (0.00078)	0.004*** (0.0008)	0.00134* (0.0007)	-0.0009 ^{ns} (0.001)	0.008 ^{ns} (0.0052)	0.011** (0.0053)
Star	0.572*** (0.02)	0.62*** (0.024)	0.47*** (0.0169)	0.464*** (0.024)	0.258*** (0.058)	0.375*** (0.059)
Experience	0.0005*** (0.0001)	0.0017 ^{ns} (0.0011)	0.001* (0.0006)	0.0011 ^{ns} (0.00097)	-0.0061 ^{ns} (0.006)	-0.00191 ^{ns} (0.003)
Designer fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Sample size	38,247	28,818	38,092	23,506	1,186	1,382

Note. *** - p-value<0.01; ** - p-value<0.05; * - p-value<0.1; ns – not significant. “Sadness” is omitted category. “Other/neutral” emotions are excluded.

As seen from Table A7, excitement (as compared to the omitted sadness) positively affects the probability of winning in all cases except for higher-tier logo contests before the AI launch. Amusement (as compared to the omitted sadness) positively affects probability of winning in lower-tier logo contests before AI, and in higher-tier logo contests before and after the AI launch. Other emotions are not significantly different from sadness in their effect on the probability of winning a contest.

Table A8. Effects of competition on emotional content and complexity

Dependent Var. / Indep. Var.	Emotions _binary (before AI)	Complexity (before AI)	Excitement (before AI)	Emotions _binary (after AI)	Complexity (after AI)	Excitement (after AI)
Number of submissions per contest	-0.0007 ^{ns} (0.00057)	-0.000016 ^{ns} (0.00001)	-0.000556 ^{ns} (0.00057)	-0.0004 ^{ns} (0.0004)	-0.000077 ^{ns} (0.000071)	-0.00041 ^{ns} (0.00038)
Number of designers per contest	-0.0064 ^{ns} (0.0051)	-0.0011 ^{ns} (0.00105)	-0.000646 ^{ns} (0.00047)	-0.0062 ^{ns} (0.00453)	-0.00145 ^{ns} (0.0052)	-0.00165 ^{ns} (0.0028)
Constant	1.526 ^{***} (0.02)	0.909 ^{***} (0.0019)	3.68 ^{***} (0.01)	1.657 ^{***} (0.02)	0.88 ^{***} (0.00178)	3.725 ^{***} (0.0098)
Designer Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
Sample Size	206,867	206,867	206,867	191,394	191,394	191,394

Note. *** - p-value<0.01; ** - p-value<0.05; * - p-value<0.1; ns – not significant.

Table A9. Effects of task requirements on emotional content and complexity

Dep. Var. Indep. Var.	Emotions _binary (before AI)	Complexity (before AI)	Excitement (before AI)	Emotions _binary (after AI)	Complexity (after AI)	Excitement (after AI)
Specific req.	-0.0012 ^{ns} (0.0011)	0.01 ^{***} (0.001)	0.017 ^{ns} (0.05)	-0.0036 ^{ns} (0.0068)	0.0041 ^{***} (0.001)	-0.0113 ^{ns} (0.057)
Abstract req.	0.0005 ^{ns} (0.00055)	0.0084 ^{***} (0.00086)	0.027 ^{ns} (0.046)	-0.0017 ^{ns} (0.005)	0.003 ^{***} (0.00088)	-0.01 ^{ns} (0.048)
Req. to convey brand emotions/feel	0.0016 ^{***} (0.0004)	0.0052 ^{***} (0.00065)	0.017 ^{***} (0.0034)	0.0006 ^{***} (0.00011)	0.00157 ^{**} (0.00067)	0.03 ^{***} (0.0036)
Constant	0.767 ^{***} (0.015)	0.887 ^{***} (0.0024)	3.55 ^{***} (0.0128)	0.775 ^{***} (0.0014)	0.859 ^{***} (0.00255)	3.687 ^{***} (0.014)
Designer Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
Sample Size	206,867	206,867	206,867	191,394	191,394	191,394

Note. *** - p-value<0.01; ** - p-value<0.05; * - p-value<0.1; ns – not significant.

Table A10. Test of parallel trends in difference-in-differences models for successful focused designers vs. control group

Dependent Variable Indep. Variable	Emotions_binary	Complexity	Excitement
Treatment_group*time_{t-6}	0.166 ^{ns} (0.48)	-0.09 ^{ns} (0.058)	0.495 ^{ns} (0.37)
Treatment_group*time_{t-5}	0.74 ^{ns} (0.507)	-0.044 ^{ns} (0.059)	0.187 ^{ns} (0.379)
Treatment_group*time_{t-4}	0.62 ^{ns} (0.51)	0.144 ^{ns} (0.111)	0.843 ^{ns} (0.657)
Treatment_group*time_{t-3}	-0.72 ^{ns} (1.19)	0.173 ^{ns} (0.242)	0.69 ^{ns} (1.33)
Treatment_group*time_{t-2}	-0.036 ^{ns} (0.63)	0.054 ^{ns} (0.069)	-0.425 ^{ns} (0.442)
Treatment_group*time_{t-1}	0.032 ^{ns} (0.6303)	0.12 ^{ns} (0.09)	-0.492 ^{ns} (0.455)
Treatment_group*time_t	Omitted Baseline	Omitted Baseline	Omitted Baseline
Constant	3.2*** (0.36)	1.79*** (0.069)	4.16*** (0.325)
Monthly Dummies	Yes	Yes	Yes
Designer Fixed Effects	Yes	Yes	Yes
Sample Size	17,823	17,123	17,823

Note. *** - p-value<0.01; ** - p-value<0.05; * - p-value<0.1; ns – not significant. Treatment_group is successful focused designers.

Table A11. Test of parallel trends in difference-in-differences models for successful cross-tier designers vs. control group

Dependent Variable Indep. Variable	Emotions_binary	Complexity	Excitement
Treatment_group*time_{t-6}	0.27 ^{ns} (0.455)	0.055 ^{ns} (0.07)	0.118 ^{ns} (0.365)
Treatment_group*time_{t-5}	0.908 ^{ns} (0.707)	0.1 ^{ns} (0.075)	0.142 ^{ns} (0.393)
Treatment_group*time_{t-4}	0.661 ^{ns} (0.497)	0.254 ^{ns} (0.173)	0.575 ^{ns} (0.458)
Treatment_group*time_{t-3}	0.302 ^{ns} (0.197)	-0.062 ^{ns} (0.052)	0.183 ^{ns} (0.27)
Treatment_group*time_{t-2}	-0.397 ^{ns} (0.609)	-0.24 ^{ns} (0.155)	-0.571 ^{ns} (0.455)
Treatment_group*time_{t-1}	-0.028 ^{ns} (0.61)	0.05 ^{ns} (0.088)	-0.322 ^{ns} (0.465)
Treatment_group*time_t	Omitted Baseline	Omitted Baseline	Omitted Baseline
Constant	3.11 ^{***} (0.32)	1.74 ^{***} (0.045)	3.708 ^{***} (0.239)
Monthly Dummies	Yes	Yes	Yes
Designer Fixed Effects	Yes	Yes	Yes
Sample Size	15,086	14,987	15,086

Note. *** - p-value<0.01; ** - p-value<0.05; * - p-value<0.1; ns – not significant. Treatment_group is successful cross-tier designers.

Table A12. Test of parallel trends in difference-in-differences models for successful cross-category designers vs. control group

Dependent Variable Indep. Variable	Emotions_binary	Complexity	Excitement
Treatment_group*time_{t-6}	0.146 ^{ns} (0.453)	-0.0695 ^{ns} (0.0665)	-0.089 ^{ns} (0.324)
Treatment_group*time_{t-5}	0.803 ^{ns} (0.647)	0.0381 ^{ns} (0.0724)	0.187 ^{ns} (0.353)
Treatment_group*time_{t-4}	0.619 ^{ns} (0.496)	0.175 ^{ns} (0.117)	0.411 ^{ns} (0.410)
Treatment_group*time_{t-3}	0.428 ^{ns} (0.978)	0.262 ^{ns} (0.273)	0.957 ^{ns} (1.65)
Treatment_group*time_{t-2}	-0.141 ^{ns} (0.608)	-0.138 ^{ns} (0.93)	-0.514 ^{ns} (0.405)
Treatment_group*time_{t-1}	0.0149 ^{ns} (0.609)	0.052 ^{ns} (0.085)	-0.379 ^{ns} (0.416)
Treatment_group*time_t	Omitted Baseline	Omitted Baseline	Omitted Baseline
Constant	3.117 ^{***} (0.323)	1.736 ^{***} (0.044)	3.708 ^{***} (0.214)
Monthly Dummies	Yes	Yes	Yes
Designer Fixed Effects	Yes	Yes	Yes
Sample Size	15,086	14,987	15,086

Note. *** - p-value<0.01; ** - p-value<0.05; * - p-value<0.1; ns – not significant. Treatment_group is successful cross-category designers.

Table A13. Comparison of the number of contests before and after the AI launch for 3 groups of successful and unsuccessful designers as compared to the control group (PSM-DID model)

Dep. Var. is the number of contests per user per day	Focused success. designers	Cross-tier success. designers	Cross- category success. designers	Focused unsucc. designers	Cross-tier unsucc. designers	Cross- category unsucc. designers
After	-0.071 ^{ns} (0.21)	0.0158 ^{ns} (0.255)	-0.004 ^{ns} (0.183)	0.07 ^{ns} (0.1)	-0.031 ^{ns} (0.119)	0.473 ^{ns} (0.333)
Treated_ group	1.578 ^{***} (0.27)	2.82 ^{***} (0.185)	1.552 ^{***} (0.189)	3.64 ^{***} (0.123)	2.67 ^{***} (0.083)	3.2 ^{***} (0.165)
DID	-0.0858 ^{ns} (0.203)	0.4 ^{ns} (0.258)	0.142 ^{ns} (0.186)	0.329 ^{**} (0.167)	0.222 [*] (0.121)	0.4 [*] (0.22)
Sample size	27,920	26,852	22,540	61,185	58,498	34,937

Note. *** - p-value<0.01; ** - p-value<0.05; * - p-value<0.1; ns – not significant

Table A14. Comparison of the number of re-submissions before and after the AI launch for 3 groups of successful and unsuccessful designers as compared to the control group (PSM-DID model)

Dep. Var. is the number of re-submiss. per user per contest	Focused success. designers	Cross-tier success. designers	Cross-category success. designers	Focused unsucc. designers	Cross-tier unsucc. designers	Cross-category unsucc. designers
After	-1.52** (0.655)	-0.092 ^{ns} (2.65)	-0.092 ^{ns} (1.95)	-0.714*** (0.231)	-0.75*** (0.112)	-0.0545 ^{ns} (0.087)
Treated_group	-2.05*** (0.753)	-3.01 ^{ns} (1.931)	-2.91** (1.419)	-1.58*** (0.159)	-4.59*** (1.26)	-0.398*** (0.063)
DID	1.121* (0.63)	3.7*** (1.19)	6.69*** (1.98)	0.3776 ^{ns} (0.237)	0.6 ^{ns} (0.437)	0.081 ^{ns} (0.0878)
Sample size	27,920	26,852	22,540	61,185	58,498	34,937

Note. *** - p-value<0.01; ** - p-value<0.05; * - p-value<0.1; ns – not significant

Table A15. Comparison of presence of emotional content before and after the AI launch for 3 groups of successful designers as compared to the control group (PSM-DID model)

Dep. variable is emotions_binary	Focused successful designers	Cross-tier successful designers	Cross-category successful designers
After	-0.18 ^{ns} (0.42)	-0.31 ^{ns} (0.37)	-0.36 ^{ns} (0.27)
Treated_group	-1.89*** (0.132)	-1.64*** (0.3)	-1.98*** (0.2)
DID	0.261** (0.129)	0.43** (0.19)	0.478** (0.22)
Constant	2.75*** (0.27)	1.3*** (0.031)	0.955** (0.467)
Sample size	17,823	15,086	14,337

Note. *** - p-value<0.01; ** - p-value<0.05; * - p-value<0.1; ns – not significant

Table A16. Comparison of complexity before and after the AI launch for 3 groups of successful designers as compared to the control group (PSM-DID model)

Dependent variable is complexity	Focused successful designers	Cross-tier successful designers	Cross-category successful designers
After	-0.054 ^{ns} (0.074)	-0.189 ^{***} (0.07)	-0.181 ^{**} (0.072)
Treated_group	-0.958 ^{***} (0.086)	-0.85 ^{***} (0.1)	-0.925 ^{***} (0.113)
DID	0.068 ^{**} (0.031)	0.197 ^{***} (0.068)	0.224 ^{***} (0.07)
Constant	1.72 ^{***} (0.047)	1.117 ^{***} (0.058)	1.779 ^{***} (0.101)
Sample size	17,123	14,987	14,149

Note. *** - p-value<0.01; ** - p-value<0.05; * - p-value<0.1; ns – not significant

Table A17. Comparison of five emotions before and after the AI launch for 3 groups of successful designers as compared to the control group (PSM-DID model)

Dependent variable is After AI	Focused successful designers	Cross-tier successful designers	Cross-category successful designers
Amusement (DID)	-0.07 ^{ns} (0.084)	0.064 ^{ns} (0.078)	0.293 ^{ns} (0.303)
Awe (DID)	0.0198 ^{ns} (0.054)	-0.144 ^{ns} (0.171)	0.051 ^{ns} (0.232)
Contentment (DID)	-0.2 ^{ns} (0.17)	0.06 ^{ns} (0.11)	0.079 ^{ns} (0.339)
Excitement (DID)	0.46 [*] (0.25)	0.81 ^{**} (0.41)	0.74 ^{**} (0.36)
Sadness (DID)	0.255 ^{ns} (0.21)	-0.2 ^{ns} (0.218)	-0.386 ^{ns} (0.758)
Sample size	17,823	15,086	14,337

Note. *** - p-value<0.01; ** - p-value<0.05; * - p-value<0.1; ns – not significant. Only DID coefficients are reported.

Bibliography

- Abrahamson, E., and Rosenkopf, L. 1993. "Institutional and Competitive Bandwagons: Using Mathematical Modeling as a Tool to Explore Innovation Diffusion," *The Academy of Management Review* (18:3), pp. 487–517. (<https://doi.org/10.2307/258906>).
- Acemoglu, D., and Restrepo, P. 2017. "Robots and Jobs: Evidence from US Labor Markets," Working Paper No. 23285, Working Paper Series, Working Paper, National Bureau of Economic Research, March. (<https://doi.org/10.3386/w23285>).
- Acemoglu, D., and Restrepo, P. 2018. "The Race between Man and Machine: Implications of Technology for Growth, Factor Shares, and Employment," *American Economic Review* (108:6), pp. 1488–1542. (<https://doi.org/10.1257/aer.20160696>).
- Adamopoulos, P., Ghose, A., and Todri, V. 2015. "Estimating the Impact of User Personality Traits on Word-of-Mouth: Text-Mining Microblogging Platforms," SSRN Scholarly Paper No. ID 2679199, SSRN Scholarly Paper, Rochester, NY: Social Science Research Network, October 23. (<https://papers.ssrn.com/abstract=2679199>).
- Adweek.com. 2015. 91% of Retail Brands Use Two or More Social Media Channels. Retrieved September 14, 2017, <https://www.adweek.com/digital/yesmail-retail-brands-social-media-channels/>
- Ahmed, S., Wallace, K. M., and Blessing, L. T. 2003. "Understanding the Differences between How Novice and Experienced Designers Approach Design Tasks," *Research in Engineering Design* (14:1), Springer, pp. 1–11.
- Aizawa, A. 2003. "An Information-Theoretic Perspective of Tf-Idf Measures," *Information Processing & Management* (39:1), pp. 45–65. ([https://doi.org/10.1016/S0306-4573\(02\)00021-3](https://doi.org/10.1016/S0306-4573(02)00021-3)).
- Akst, D. 2013. "Automation Anxiety," *The Wilson Quarterly* (37:3), Woodrow Wilson International Center for Scholars, p. 65.
- Allison, P. D., Williams, R., and Moral-Benito, E. 2017. "Maximum Likelihood for Cross-Lagged Panel Models with Fixed Effects," *Socius* (3), SAGE Publications Sage CA: Los Angeles, CA, p. 2378023117710578.
- Araújo, R. M. 2013. "99designs: An Analysis of Creative Competition in Crowdsourced Design," in *HCOMP*.

- Archak, N. 2010. "Money, Glory and Cheap Talk: Analyzing Strategic Behavior of Contestants in Simultaneous Crowdsourcing Contests on TopCoder.Com," in *Proceedings of the 19th International Conference on World Wide Web, WWW '10*, Raleigh, North Carolina, USA: Association for Computing Machinery, April 26, pp. 21–30. (<https://doi.org/10.1145/1772690.1772694>).
- Archak, N., and Sundararajan, A. 2009. "Optimal Design of Crowdsourcing Contests," *ICIS 2009 Proceedings*. (<https://aisel.aisnet.org/icis2009/200>).
- Ashley, C., and Tuten, T. 2015. "Creative Strategies in Social Media Marketing: An Exploratory Study of Branded Social Content and Consumer Engagement," *Psychology & Marketing* (32:1), pp. 15–27. (<https://doi.org/10.1002/mar.20761>).
- Athar, S., and Wang, Z. 2019. "A Comprehensive Performance Evaluation of Image Quality Assessment Algorithms," *IEEE Access* (7), pp. 140030–140070. (<https://doi.org/10.1109/ACCESS.2019.2943319>).
- Atman, C. J., Chimka, J. R., Bursic, K. M., and Nachtmann, H. L. 1999. "A Comparison of Freshman and Senior Engineering Design Processes," *Design Studies* (20:2), pp. 131–152. ([https://doi.org/10.1016/S0142-694X\(98\)00031-3](https://doi.org/10.1016/S0142-694X(98)00031-3)).
- Aulakh, P. S., Rotate, M., and Teegen, H. 2000. "Export Strategies and Performance of Firms from Emerging Economies: Evidence from Brazil, Chile, and Mexico," *Academy of Management Journal* (43:3), pp. 342–361. (<https://doi.org/10.5465/1556399>).
- Autio, E., Nambisan, S., Thomas, L. D. W., and Wright, M. 2018. "Digital Affordances, Spatial Affordances, and the Genesis of Entrepreneurial Ecosystems," *Strategic Entrepreneurship Journal* (12:1), pp. 72–95. (<https://doi.org/10.1002/sej.1266>).
- Autor, D. H. 2003. "Outsourcing at Will: The Contribution of Unjust Dismissal Doctrine to the Growth of Employment Outsourcing," *Journal of Labor Economics* (21:1), The University of Chicago Press, pp. 1–42. (<https://doi.org/10.1086/344122>).
- Autor, D. H. 2015. "Why Are There Still So Many Jobs? The History and Future of Workplace Automation," *Journal of Economic Perspectives* (29:3), pp. 3–30. (<https://doi.org/10.1257/jep.29.3.3>).
- Autor, D., Dorn, D., Katz, L. F., Patterson, C., and Van Reenen, J. 2017. "The Fall of the Labor Share and the Rise of Superstar Firms," Working Paper No. 23396, Working Paper, National Bureau of Economic Research, May. (<https://doi.org/10.3386/w23396>).

- Avital, M., and Te'eni, D. 2009. "From Generative Fit to Generative Capacity: Exploring an Emerging Dimension of Information Systems Design and Task Performance," *Information Systems Journal* (19:4), pp. 345–367. (<https://doi.org/10.1111/j.1365-2575.2007.00291.x>).
- Badir, Y., Charoenngam, C., Siengthai, S., and Teeratansirikool, L. 2013. "Competitive Strategies and Firm Performance: The Mediating Role of Performance Measurement," *International Journal of Productivity and Performance Management* (62:2), pp. 168–184. (<https://doi.org/10.1108/17410401311295722>).
- Bajaj, A., and Bond, S. D. 2018. "Beyond Beauty: Design Symmetry and Brand Personality," *Journal of Consumer Psychology* (28:1), pp. 77–98. (<https://doi.org/10.1002/jcpy.1009>).
- Bapna, S., Benner, M., and Qiu, L. 2019. "Nurturing Online Communities: An Empirical Investigation," *Management Information Systems Quarterly* (43:2), pp. 425–452.
- Barney, J. 1991. "Firm Resources and Sustained Competitive Advantage," *Journal of Management* (17:1), pp. 99–120. (<https://doi.org/10.1177/014920639101700108>).
- Beckert, J. 2010. "Institutional Isomorphism Revisited: Convergence and Divergence in Institutional Change*," *Sociological Theory* (28:2), pp. 150–166. (<https://doi.org/10.1111/j.1467-9558.2010.01369.x>).
- Bergen, M., and Peteraf, M. A. 2002. "Competitor Identification and Competitor Analysis: A Broad-Based Managerial Approach," *Managerial and Decision Economics* (23:4–5), pp. 157–169. (<https://doi.org/10.1002/mde.1059>).
- Berlyne, D. E. 1974. *Studies in the New Experimental Aesthetics: Steps toward an Objective Psychology of Aesthetic Appreciation.*, Hemisphere.
- Berrett, Ti., and Slack, T. 1999. "An Analysis of the Influence of Competitive and Institutional Pressures on Corporate Sponsorship Decisions," *Journal of Sport Management* (13:2), p. 114. (<https://doi.org/10.1123/jsm.13.2.114>).
- Björklund, T. A. 2013. "Initial Mental Representations of Design Problems: Differences between Experts and Novices," *Design Studies* (34:2), pp. 135–160. (<https://doi.org/10.1016/j.destud.2012.08.005>).
- Blog.depict.com. 2015. Neiman Marcus Make Some Noice. Retrieved September 14, 2017, <http://blog.depict.com/post/113531337225/were-excited-to-see-neiman-marcus-make-some?route=%2Fpost%2F%3Aid%2F%3Asummary>

- Blog.hubspot.com. 2019. How to Use Hashtags on Twitter: A Simple Guide for Marketers. Retrieved February 14, 2019, <https://blog.hubspot.com/blog/tabid/6307/bid/32497/how-to-use-hashtags-on-twitter-a-simple-guide-for-marketers.aspx>
- Bockstedt, J., Druehl, C., and Mishra, A. 2016. "Heterogeneous Submission Behavior and Its Implications for Success in Innovation Contests with Public Submissions," *Production and Operations Management* (25:7), pp. 1157–1176. (<https://doi.org/10.1111/poms.12552>).
- Borth, D., Ji, R., Chen, T., Breuel, T., and Chang, S.-F. 2013. "Large-Scale Visual Sentiment Ontology and Detectors Using Adjective Noun Pairs," in *Proceedings of the 21st ACM International Conference on Multimedia*, MM '13, New York, NY, USA: ACM, pp. 223–232. (<https://doi.org/10.1145/2502081.2502282>).
- Boudreau, K., Helfat, C. E., Lakhani, K. R., and Menietti, M. E. 2012. "Field Evidence on Individual Behavior & Performance in Rank-Order Tournaments," *Harvard Business School Working Paper Series# 13-016*.
- Braga, A., and Logan, R. K. 2017. "The Emperor of Strong AI Has No Clothes: Limits to Artificial Intelligence," *Information* (8:4), Multidisciplinary Digital Publishing Institute, p. 156. (<https://doi.org/10.3390/info8040156>).
- Brougham, D., and Haar, J. 2018. "Smart Technology, Artificial Intelligence, Robotics, and Algorithms (STARA): Employees' Perceptions of Our Future Workplace," *Journal of Management & Organization* (24:2), Cambridge University Press, pp. 239–257. (<https://doi.org/10.1017/jmo.2016.55>).
- Brouthers, L. E., O'Donnell, E., and Hadjimarcou, J. 2005. "Generic Product Strategies for Emerging Market Exports into Triad Nation Markets: A Mimetic Isomorphism Approach," *Journal of Management Studies* (42:1), pp. 225–245. (<https://doi.org/10.1111/j.1467-6486.2005.00495.x>).
- Bruns, A., Woodford, D., and Sadkowsky, T. 2014. "Towards a Methodology for Examining Twitter Follower Accession," *First Monday* (19:4). (<http://firstmonday.org/ojs/index.php/fm/article/view/5211>).
- Brynjolfsson, E., and McAfee, A. 2012. "Winning the Race with Ever-Smarter Machines," *MIT Sloan Management Review* (53:2), Massachusetts Institute of Technology, Cambridge, MA, p. 53.
- Brynjolfsson, E., and McAfee, A. 2014. *The Second Machine Age: Work, Progress, and Prosperity in a Time of Brilliant Technologies*, WW Norton & Company.

- Campaignlive.co.uk. 2019. Twitter picks the most creative brand campaigns of 2018. Retrieved May 10, 2019, <https://www.campaignlive.co.uk/article/twitter-picks-creative-brand-campaigns-2018/1524021>
- Castaldi, C., Frenken, K., and Los, B. 2015. "Related Variety, Unrelated Variety and Technological Breakthroughs: An Analysis of US State-Level Patenting," *Regional Studies* (49:5), Routledge, pp. 767–781. (<https://doi.org/10.1080/00343404.2014.940305>).
- Chen, C., Young, D., and Zhuang, Z. 2012. "Externalities of Mandatory IFRS Adoption: Evidence from Cross-Border Spillover Effects of Financial Information on Investment Efficiency," *The Accounting Review* (88:3), American Accounting Association, pp. 881–914. (<https://doi.org/10.2308/accr-50384>).
- Chen, P.-Y., Pavlou, P. A., and Yang, Y. 2014. "Determinants of Open Contest Participation in Online Labor Markets," SSRN Scholarly Paper No. ID 2510114, SSRN Scholarly Paper, Rochester, NY: Social Science Research Network, September 1. (<https://doi.org/10.2139/ssrn.2510114>).
- Cheng, Z., and Nault, B. R. 2012. "Relative Industry Concentration and Customer-Driven IT Spillovers," *Information Systems Research* (23:2), INFORMS, pp. 340–355.
- Choi, J. P., Jeon, D.-S., and Kim, B.-C. 2019. "Privacy and Personal Data Collection with Information Externalities," *Journal of Public Economics* (173), pp. 113–124. (<https://doi.org/10.1016/j.jpubeco.2019.02.001>).
- Chui, M., Manyika, J., and Miremadi, M. 2016. "Where Machines Could Replace Humans—and Where They Can't (Yet)," *McKinsey Quarterly* (30:2), pp. 1–9.
- Ciocca, G., Corchs, S., Gasparini, F., Bricolo, E., and Tebano, R. 2015. "Does Color Influence Image Complexity Perception?," in *Computational Color Imaging*, Lecture Notes in Computer Science, A. Trémeau, R. Schettini, and S. Tominaga (eds.), Cham: Springer International Publishing, pp. 139–148. (https://doi.org/10.1007/978-3-319-15979-9_14).
- Constantinides, P., Henfridsson, O., and Parker, G. G. 2018. "Introduction—Platforms and Infrastructures in the Digital Age," *Information Systems Research* (29:2), INFORMS, pp. 381–400. (<https://doi.org/10.1287/isre.2018.0794>).
- Correa, T., Hinsley, A. W., and de Zúñiga, H. G. 2010. "Who Interacts on the Web?: The Intersection of Users' Personality and Social Media Use," *Computers in Human Behavior* (26:2), pp. 247–253. (<https://doi.org/10.1016/j.chb.2009.09.003>).

- Cross, N. 2004. "Expertise in Design: An Overview," *Design Studies* (25:5), Expertise in Design, pp. 427–441. (<https://doi.org/10.1016/j.destud.2004.06.002>).
- Culnan, M. J., McHugh, P. J., and Zubillaga, J. I. 2010. "How Large U.s. Companies Can Use Twitter and Other Social Media to Gain Business Value," *MIS Quarterly Executive* (9:4), pp. 243–259.
- Dailymail.co.uk. 2014. How Puma's new ad campaign featuring world's most famous sports stars backfired spectacularly after it invited Twitter users to get involved. Retrieved May 30, 2019, <https://www.dailymail.co.uk/news/article-2728786/Pumas-ForeverFaster-ad-campaign-backfires-global-Twitter-joke.html>
- De Marchis, G. P., Reales-Avilés, J. M., and del Prado Rivero, M. 2018. "Comparative Values of Variables Related to Brand Logos," *Measuring Business Excellence*, Emerald Publishing Limited.
- Deck, C., and Kimbrough, E. O. 2017. "Experimenting with Contests for Experimentation," *Southern Economic Journal* (84:2), pp. 391–406. (<https://doi.org/10.1002/soej.12185>).
- Dessart, L., Veloutsou, C., and Morgan-Thomas, A. 2015. "Consumer Engagement in Online Brand Communities: A Social Media Perspective," *Journal of Product & Brand Management* (24:1), pp. 28–42. (<https://doi.org/10.1108/JPBM-06-2014-0635>).
- DiMaggio, P. J., and Powell, W. W. 1983. "The Iron Cage Revisited: Institutional Isomorphism and Collective Rationality in Organizational Fields," *American Sociological Review* (48:2), pp. 147–160. (<https://doi.org/10.2307/2095101>).
- DiPalantino, D., Karagiannis, T., and Vojnovic, M. 2011. *Individual and Collective User Behavior in Crowdsourcing Services*.
- DiPalantino, D., and Vojnovic, M. 2009. "Crowdsourcing and All-Pay Auctions," in *Proceedings of the 10th ACM Conference on Electronic Commerce, EC '09*, Stanford, California, USA: Association for Computing Machinery, July 6, pp. 119–128. (<https://doi.org/10.1145/1566374.1566392>).
- Dixon, J., Hong, B., and Wu, L. 2019. "The Employment Consequences of Robots: Firm-Level Evidence," SSRN Scholarly Paper No. ID 3422581, SSRN Scholarly Paper, Rochester, NY: Social Science Research Network, July 19. (<https://doi.org/10.2139/ssrn.3422581>).
- Duel.tech. 2018. 7 Effective Examples of Brand Community Hashtags. Retrieved December 24, 2019, <https://www.duel.tech/blog/10-effective-examples-brand-community-hashtags>

- Eberhard, B., Podio, M., Alonso, A. P., Radovica, E., Avotina, L., Peiseniece, L., Caamaño Sendon, M., Gonzales Lozano, A., and Solé-Pla, J. 2017. "Smart Work: The Transformation of the Labour Market Due to the Fourth Industrial Revolution (I4. 0).," *International Journal of Business & Economic Sciences Applied Research* (10:3).
- Einwiller, S. A., and Steilen, S. 2015. "Handling Complaints on Social Network Sites – An Analysis of Complaints and Complaint Responses on Facebook and Twitter Pages of Large US Companies," *Public Relations Review* (41:2), Digital Publics, pp. 195–204. (<https://doi.org/10.1016/j.pubrev.2014.11.012>).
- Elgammal, A., Liu, B., Elhoseiny, M., and Mazzone, M. 2017. "CAN: Creative Adversarial Networks, Generating 'Art' by Learning About Styles and Deviating from Style Norms," *ArXiv:1706.07068 [Cs]*. (<http://arxiv.org/abs/1706.07068>).
- Ericsson, K. A., Hoffman, R. R., and Kozbelt, A. 2018. *The Cambridge Handbook of Expertise and Expert Performance*, Cambridge University Press.
- Farndale, E., and Paauwe, J. 2007. "Uncovering Competitive and Institutional Drivers of HRM Practices in Multinational Corporations," *Human Resource Management Journal* (17:4), pp. 355–375. (<https://doi.org/10.1111/j.1748-8583.2007.00050.x>).
- Fleming, L. 2001. "Recombinant Uncertainty in Technological Search," *Management Science* (47:1), INFORMS, pp. 117–132. (<https://doi.org/10.1287/mnsc.47.1.117.10671>).
- Frenken, K., Izquierdo, L. R., and Zeppini, P. 2012. "Branching Innovation, Recombinant Innovation, and Endogenous Technological Transitions," *Environmental Innovation and Societal Transitions* (4), pp. 25–35. (<https://doi.org/10.1016/j.eist.2012.06.001>).
- Frey, C. B., and Osborne, M. A. 2017. "The Future of Employment: How Susceptible Are Jobs to Computerisation?," *Technological Forecasting and Social Change* (114), pp. 254–280. (<https://doi.org/10.1016/j.techfore.2016.08.019>).
- Gao, X. A., Bachrach, Y., Key, P. B., and Graepel, T. 2012. "Quality Expectation-Variance Tradeoffs in Crowdsourcing Contests," in *AAAI*.
- Gatys, L. A., Ecker, A. S., and Bethge, M. 2015. "A Neural Algorithm of Artistic Style," *ArXiv:1508.06576 [Cs, q-Bio]*. (<http://arxiv.org/abs/1508.06576>).
- Gavilanes, J. M., Flatten, T. C., and Brettel, M. 2018. "Content Strategies for Digital Consumer Engagement in Social Networks: Why Advertising Is an Antecedent of Engagement," *Journal of Advertising* (47:1), pp. 4–23. (<https://doi.org/10.1080/00913367.2017.1405751>).

- Ghose, A., Ipeirotis, P., and Li, B. 2012. "Designing Ranking Systems for Hotels on Travel Search Engines by Mining User-Generated and Crowdsourced Content," *Marketing Science*, pp. 493–520. (<https://doi.org/10.1287/mksc.1110.0700>).
- Ghosh, A., and Kleinberg, R. 2016. *Optimal Contest Design for Simple Agents*, Association for Computing Machinery. (<https://doi.org/10.1145/2930955>).
- Gibson, J. J. 1986. *The Ecological Approach to Visual Perception*, Lawrence Erlbaum Associates.
- Goh, K.-Y., Heng, C.-S., and Lin, Z. 2013. "Social Media Brand Community and Consumer Behavior: Quantifying the Relative Impact of User- and Marketer-Generated Content," *Information Systems Research* (24:1), pp. 88–107. (<https://doi.org/10.1287/isre.1120.0469>).
- Goodman, L. A. 1961. "Snowball Sampling," *The Annals of Mathematical Statistics* (32:1), pp. 148–170.
- Graetz, G., and Michaels, G. 2018. "Robots at Work," *Review of Economics and Statistics* (100:5), MIT Press, pp. 753–768.
- Grinsven, B. van, and Das, E. 2016. "Logo Design in Marketing Communications: Brand Logo Complexity Moderates Exposure Effects on Brand Recognition and Brand Attitude," *Journal of Marketing Communications* (22:3), Routledge, pp. 256–270. (<https://doi.org/10.1080/13527266.2013.866593>).
- Gross, D. P. 2016. "Creativity under Fire: The Effects of Competition on Creative Production," *Review of Economics and Statistics*, MIT Press, pp. 1–49.
- Guillén, M. F. 2003. *The Limits of Convergence: Globalization and Organizational Change in Argentina, South Korea, and Spain*, Princeton University Press.
- Gur, F. A., and Greckhamer, T. 2018. "Know Thy Enemy: A Review and Agenda for Research on Competitor Identification," *Journal of Management*, p. 0149206317744250. (<https://doi.org/10.1177/0149206317744250>).
- Hawley, A. H. 1968. "Human Ecology//International Encyclopedia of Social Sciences. L," NY: *Manncmillan*, pp. 328–337.
- He, S., Rui, H., and Whinston, A. B. 2017. "Social Media Strategies in Product-Harm Crises," *Information Systems Research*. (<https://doi.org/10.1287/isre.2017.0707>).
- He, X., and Zhang, W. 2018. "Emotion Recognition by Assisted Learning with Convolutional Neural Networks," *Neurocomputing* (291), pp. 187–194. (<https://doi.org/10.1016/j.neucom.2018.02.073>).

- Helms, R. W., Booij, E., and Spruit, M. R. 2012. "Reaching Out: Involving Users in Innovation Tasks Through Social Media," *ECIS 2012 Proceedings*. (<https://aisel.aisnet.org/ecis2012/193>).
- Hertzmann, A. 2018. "Can Computers Create Art?," *Arts* (7:2), Multidisciplinary Digital Publishing Institute, p. 18. (<https://doi.org/10.3390/arts7020018>).
- Hitt, L., and Tambe, P. 2006. "Measuring Spillovers From Information Technology Investments," *ICIS 2006 Proceedings*. (<https://aisel.aisnet.org/icis2006/108>).
- Hoberg, G., and Phillips, G. 2016. "Text-Based Network Industries and Endogenous Product Differentiation," *Journal of Political Economy* (124:5), pp. 1423–1465. (<https://doi.org/10.1086/688176>).
- Hofstetter, R., Nair, H., and Misra, S. 2020. "Can Open Innovation Survive? Imitation and Return on Originality in Crowdsourcing Creative Work," SSRN Scholarly Paper No. ID 3133158, SSRN Scholarly Paper, Rochester, NY: Social Science Research Network, January 14. (<https://papers.ssrn.com/abstract=3133158>).
- Huang, A. 2008. "Similarity Measures for Text Document Clustering," in *Proceedings of the Sixth New Zealand Computer Science Research Student Conference (NZCSRSC2008)*, Christchurch, New Zealand, pp. 49–56.
- Huang, M.-H., Rust, R., and Maksimovic, V. 2019. "The Feeling Economy: Managing in the Next Generation of Artificial Intelligence (AI)," *California Management Review* (61:4), SAGE Publications Inc, pp. 43–65. (<https://doi.org/10.1177/0008125619863436>).
- Hutto, C. J., and Gilbert, E. 2014. "Vader: A Parsimonious Rule-Based Model for Sentiment Analysis of Social Media Text," in *Eighth International AAAI Conference on Weblogs and Social Media*.
- Ibrahim, N. F., Wang, X., and Bourne, H. 2017. "Exploring the Effect of User Engagement in Online Brand Communities: Evidence from Twitter," *Computers in Human Behavior* (72), pp. 321–338. (<https://doi.org/10.1016/j.chb.2017.03.005>).
- Janze, C. 2016. "Spillover Effects in User-generated Content: Evidence from Online Reviews of Interdependent Service Chains," *PACIS 2016 Proceedings*. (<https://aisel.aisnet.org/pacis2016/192>).
- Jian, L., Li, Z., and Liu, T. X. 2017. "Simultaneous versus Sequential All-Pay Auctions: An Experimental Study," *Experimental Economics* (20:3), pp. 648–669. (<https://doi.org/10.1007/s10683-016-9504-1>).
- Jiang, Z. (Zoey), Huang, Y., and Beil, D. R. 2018. "The Role of Feedback in Dynamic Crowdsourcing Contests: A Structural Empirical Analysis," SSRN Scholarly

Paper No. ID 2884922, SSRN Scholarly Paper, Rochester, NY: Social Science Research Network, December 6. (<https://doi.org/10.2139/ssrn.2884922>).

Jiang, Z. (Zoey), Huang, Y., and Beil, D. R. 2019. "The Role of Problem Specification in Crowdsourcing Contests for Design Problems: A Theoretical and Empirical Analysis," SSRN Scholarly Paper No. ID 3260862, SSRN Scholarly Paper, Rochester, NY: Social Science Research Network, October 16. (<https://doi.org/10.2139/ssrn.3260862>).

Jing, B. 2018. *Showrooming and Webrooming: Information Externalities Between Online and Offline Sellers*, INFORMS. (<https://doi.org/10.1287/mksc.2018.1084>).

Karahanna, E., Xin Xu, S., Xu, Y., and Zhang, N. (Andy). 2018. "The Needs–Affordances–Features Perspective for the Use of Social Media," *MIS Quarterly* (42), pp. 737–756. (<https://doi.org/10.25300/misq/2018/11492>).

Karras, T., Laine, S., and Aila, T. 2019. "A Style-Based Generator Architecture for Generative Adversarial Networks," *ArXiv:1812.04948 [Cs, Stat]*. (<http://arxiv.org/abs/1812.04948>).

Khasraghi, H. J., and Aghaie, A. 2014. "Crowdsourcing Contests: Understanding the Effect of Competitors' Participation History on Their Performance," *Behaviour & Information Technology* (33:12), Taylor & Francis, pp. 1383–1395. (<https://doi.org/10.1080/0144929X.2014.883551>).

Kietzmann, J. H., Hermkens, K., McCarthy, I. P., and Silvestre, B. S. 2011. "Social Media? Get Serious! Understanding the Functional Building Blocks of Social Media," *Business Horizons* (54:3), SPECIAL ISSUE: SOCIAL MEDIA, pp. 241–251. (<https://doi.org/10.1016/j.bushor.2011.01.005>).

Kim, J., Lee, M., Cho, D., and Lee, B. 2016. "Are All Spillovers Created Equal? The Impact of Blockbusters and the Composition of Backers in Online Crowdfunding," *ICIS 2016 Proceedings*. (<https://aisel.aisnet.org/icis2016/Crowdsourcing/Presentations/12>).

Kim, K., and Viswanathan, S. 2018. "The 'experts' in the Crowd: The Role of Experienced Investors in a Crowdfunding Market," *Mis Quarterly*.

Kim, M. J., and Lim, J. H. 2019. "A Comprehensive Review on Logo Literature: Research Topics, Findings, and Future Directions," *Journal of Marketing Management* (35:13–14), Routledge, pp. 1291–1365. (<https://doi.org/10.1080/0267257X.2019.1604563>).

Körpeoğlu, E., and Cho, S.-H. 2017. "Incentives in Contests with Heterogeneous Solvers," *Management Science* (64:6), INFORMS, pp. 2709–2715. (<https://doi.org/10.1287/mnsc.2017.2738>).

- Körpeoğlu, E., Korpeoglu, C. G., and Hafalir, I. E. 2017. "Parallel Innovation Contests," *Available at SSRN 2924817*.
- Krijestorac, H., Garg, R., Mahajan, V., and Hofstede, F. ter. 2017. "Understanding Spillover Effects in Consumption of Rich Digital Media," *ICIS 2017 Proceedings*. (<https://aisel.aisnet.org/icis2017/SocialMedia/Presentations/13>).
- Kumar, V., Bhaskaran, V., Mirchandani, R., and Shah, M. 2013. "Practice Prize Winner—Creating a Measurable Social Media Marketing Strategy: Increasing the Value and ROI of Intangibles and Tangibles for Hokey Pokey," *Marketing Science* (32:2), INFORMS, pp. 194–212. (<https://doi.org/10.1287/mksc.1120.0768>).
- Kwak, H., Lee, C., Park, H., and Moon, S. 2010. "What Is Twitter, a Social Network or a News Media?," in *Proceedings of the 19th International Conference on World Wide Web, WWW '10*, New York, NY, USA: ACM, pp. 591–600. (<https://doi.org/10.1145/1772690.1772751>).
- Kwark, Y., Lee, G. M., Pavlou, P., and Qiu, L. 2016. "The Spillover Effects of User-Generated Online Product Reviews on Purchases: Evidence from Clickstream Data," *ICIS 2016 Proceedings*. (<https://aisel.aisnet.org/icis2016/EBusiness/Presentations/11>).
- LaToza, T. D., Chen, M., Jiang, L., Zhao, M., and Hoek, A. v d. 2015. "Borrowing from the Crowd: A Study of Recombination in Software Design Competitions," in *2015 IEEE/ACM 37th IEEE International Conference on Software Engineering* (Vol. 1), May, pp. 551–562. (<https://doi.org/10.1109/ICSE.2015.72>).
- Lee, D., Hosanagar, K., and Nair, H. S. 2018. "Advertising Content and Consumer Engagement on Social Media: Evidence from Facebook," *Management Science* (64:11), pp. 5105–5131. (<https://doi.org/10.1287/mnsc.2017.2902>).
- Leszczensky, L., and Wolbring, T. 2019. "How to Deal with Reverse Causality Using Panel Data? Recommendations for Researchers Based on a Simulation Study," *Sociological Methods & Research*, Sage Publications Sage CA: Los Angeles, CA, p. 0049124119882473.
- Li, J. (Justin), Bonn, M. A., and Ye, B. H. 2019. "Hotel Employee's Artificial Intelligence and Robotics Awareness and Its Impact on Turnover Intention: The Moderating Roles of Perceived Organizational Support and Competitive Psychological Climate," *Tourism Management* (73), pp. 172–181. (<https://doi.org/10.1016/j.tourman.2019.02.006>).
- Liu, X., Li, N., and Xia, Y. 2019. "Affective Image Classification by Jointly Using Interpretable Art Features and Semantic Annotations," *Journal of Visual*

- Communication and Image Representation* (58), pp. 576–588. (<https://doi.org/10.1016/j.jvcir.2018.12.032>).
- Liu, X., and Lynch, L. 2011. “Do Agricultural Land Preservation Programs Reduce Farmland Loss? Evidence from a Propensity Score Matching Estimator,” *Land Economics* (87:2), [Board of Regents of the University of Wisconsin System, University of Wisconsin Press], pp. 183–201.
- Lonelybrand.com. 2017. The Gap’s Styld.by Campaign Uses Bloggers to Build Connections. Retrieved December 24, 2019, <https://lonelybrand.com/blog/the-gaps-styld-by-campaign-uses-bloggers-to-build-connections/>
- Lovejoy, K., and Saxton, G. D. 2012. “Information, Community, and Action: How Nonprofit Organizations Use Social Media*,” *Journal of Computer-Mediated Communication* (17:3), pp. 337–353. (<https://doi.org/10.1111/j.1083-6101.2012.01576.x>).
- Lu, V. N., Wirtz, J., Kunz, W. H., Paluch, S., Gruber, T., Martins, A., and Patterson, P. G. 2020. “Service Robots, Customers and Service Employees: What Can We Learn from the Academic Literature and Where Are the Gaps?,” *Journal of Service Theory and Practice* (ahead-of-print:ahead-of-print). (<https://doi.org/10.1108/JSTP-04-2019-0088>).
- Luarn, P., Lin, Y.-F., and Chiu, Y.-P. 2015. “Influence of Facebook Brand-Page Posts on Online Engagement,” *Online Information Review* (39:4), pp. 505–519.
- Luca, M. 2015. “Chapter 12 - User-Generated Content and Social Media,” in *Handbook of Media Economics* (Vol. 1), Handbook of Media Economics, S. P. Anderson, J. Waldfogel, and D. Strömberg (eds.), North-Holland, pp. 563–592. (<https://doi.org/10.1016/B978-0-444-63685-0.00012-7>).
- Majchrzak, A., Faraj, S., Kane, G. C., and Azad, B. 2013. “The Contradictory Influence of Social Media Affordances on Online Communal Knowledge Sharing,” *Journal of Computer-Mediated Communication* (19:1), pp. 38–55. (<https://doi.org/10.1111/jcc4.12030>).
- Makridakis, S. 2017. “The Forthcoming Artificial Intelligence (AI) Revolution: Its Impact on Society and Firms,” *Futures* (90), pp. 46–60. (<https://doi.org/10.1016/j.futures.2017.03.006>).
- Malhotra, A., Malhotra, C. K., and See, A. 2012. “How to Get Your Messages Retweeted,” *MIT Sloan Management Review* (53:2), pp. 61–66.
- Malhotra, A., Malhotra, C. K., and See, A. 2013. “How to Create Brand Engagement on Facebook,” *MIT Sloan Management Review* (54:2), pp. 18–20.

- Mandviwalla, M., and Watson, R. 2014. "Generating Capital from Social Media," *MIS Quarterly Executive* (13:2), pp. 97–113.
- Mann, K., and Püttmann, L. 2018. "Benign Effects of Automation: New Evidence from Patent Texts," SSRN Scholarly Paper No. ID 2959584, SSRN Scholarly Paper, Rochester, NY: Social Science Research Network, August 15. (<https://doi.org/10.2139/ssrn.2959584>).
- Manyika, J., Chui, M., and Miremadi, M. 2017. "A Future That Works: AI, Automation, Employment, and Productivity," *McKinsey Global Institute Research, Tech. Rep* (60).
- March, J. G. 1991. "Exploration and Exploitation in Organizational Learning," *Organization Science* (2:1), INFORMS, pp. 71–87.
- Marin, M. M., Lampatz, A., Wandl, M., and Leder, H. 2016. "Berlyne Revisited: Evidence for the Multifaceted Nature of Hedonic Tone in the Appreciation of Paintings and Music," *Frontiers in Human Neuroscience* (10). (<https://doi.org/10.3389/fnhum.2016.00536>).
- Mazzone, M., and Elgammal, A. 2019. "Art, Creativity, and the Potential of Artificial Intelligence," *Arts* (8:1), Multidisciplinary Digital Publishing Institute, p. 26. (<https://doi.org/10.3390/arts8010026>).
- Meeder, B., Karrer, B., Sayedi, A., Ravi, R., Borgs, C., and Chayes, J. 2011. "We Know Who You Followed Last Summer: Inferring Social Link Creation Times in Twitter," in *Proceedings of the 20th International Conference on World Wide Web*, WWW '11, New York, NY, USA: ACM, pp. 517–526. (<https://doi.org/10.1145/1963405.1963479>).
- Menon, N. 2018. "Information Spillover and Semi-Collaborative Networks in Insurer Fraud Detection," *Management Information Systems Quarterly* (42:2), pp. 407–426.
- Meyer, J. W., and Rowan, B. 1977. "Institutionalized Organizations: Formal Structure as Myth and Ceremony," *American Journal of Sociology* (83:2), pp. 340–363.
- Mount, M., and Garcia Martinez, M. 2014. "Social Media: a Tool for Open Innovation," *California Management Review* (56:4), pp. 124–143. (<https://doi.org/10.1525/cmr.2014.56.4.124>).
- Nam, H., Joshi, Y. V., and Kannan, P. K. 2017. "Harvesting Brand Information from Social Tags:," *Journal of Marketing*. (<https://doi.org/10.1509/jm.16.0044>).
- Nam, T. 2019. "Technology Usage, Expected Job Sustainability, and Perceived Job Insecurity," *Technological Forecasting and Social Change* (138), pp. 155–165. (<https://doi.org/10.1016/j.techfore.2018.08.017>).

- Novak, R., and Nikulin, Y. 2016. "Improving the Neural Algorithm of Artistic Style," *ArXiv:1605.04603 [Cs]*. (<http://arxiv.org/abs/1605.04603>).
- Pan, Y., Huang, P., and Gopal, A. 2015. "Enemies (Almost) at the Gate: New Entry Threats and R&D Investments in the U.S. IT Industry," SSRN Scholarly Paper No. ID 2607674, SSRN Scholarly Paper, Rochester, NY: Social Science Research Network, May 18. (<https://papers.ssrn.com/abstract=2607674>).
- Pant, G., and Sheng, O. R. L. 2015. "Web Footprints of Firms: Using Online Isomorphism for Competitor Identification," *Information Systems Research* (26:1), pp. 188–209. (<https://doi.org/10.1287/isre.2014.0563>).
- Parsons, A. 2013. "Using Social Media to Reach Consumers: A Content Analysis of Official Facebook Pages.," *Academy of Marketing Studies Journal* (17:2).
- Pexman, P. M., Heard, A., Lloyd, E., and Yap, M. J. 2017. "The Calgary Semantic Decision Project: Concrete/Abstract Decision Data for 10,000 English Words," *Behavior Research Methods* (49:2), pp. 407–417. (<https://doi.org/10.3758/s13428-016-0720-6>).
- Pieters, R., Wedel, M., and Batra, R. 2010. "The Stopping Power of Advertising: Measures and Effects of Visual Complexity," *Journal of Marketing* (74:5), SAGE Publications Inc, pp. 48–60. (<https://doi.org/10.1509/jmkg.74.5.048>).
- Pletikosa Cvijikj, I., and Michahelles, F. 2013. "Online Engagement Factors on Facebook Brand Pages," *Social Network Analysis and Mining* (3:4), pp. 843–861. (<https://doi.org/10.1007/s13278-013-0098-8>).
- Popovic, V. 2004. "Expertise Development in Product Design—Strategic and Domain-Specific Knowledge Connections," *Design Studies* (25:5), Expertise in Design, pp. 527–545. (<https://doi.org/10.1016/j.destud.2004.05.006>).
- Pozzi, G., Pigni, F., and Vitari, C. 2014. "Affordance Theory in the IS Discipline: A Review and Synthesis of the Literature," in *AMCIS*.
- Rap-up.com. 2016. Rihanna to Launch First Fenty X Puma Collection in September. Retrieved September 14, 2017, <https://www.rap-up.com/2016/08/22/rihanna-fenty-x-puma-collection-launch-date/>
- Rauschnabel, P. A., Sheldon, P., and Herzfeldt, E. 2019. "What Motivates Users to Hashtag on Social Media?," *Psychology & Marketing* (36:5), pp. 473–488. (<https://doi.org/10.1002/mar.21191>).
- de Reuver, M., Sørensen, C., and Basole, R. C. 2018. "The Digital Platform: A Research Agenda," *Journal of Information Technology* (33:2), SAGE Publications Ltd, pp. 124–135. (<https://doi.org/10.1057/s41265-016-0033-3>).

- Riedl, C., and Seidel, V. P. 2018. "Learning from Mixed Signals in Online Innovation Communities," *Organization Science* (29:6), INFORMS, pp. 1010–1032. (<https://doi.org/10.1287/orsc.2018.1219>).
- Roberts, D. L., and Piller, F. T. 2016. "Finding the Right Role for Social Media in Innovation," *MIT Sloan Management Review* (57:3), p. 41.
- Rosenbaum, P. R., and Rubin, D. B. 1983. "The Central Role of the Propensity Score in Observational Studies for Causal Effects," *Biometrika* (70:1), Oxford University Press, pp. 41–55.
- Ryan, A. M., Burgess, J. F., and Dimick, J. B. 2015. "Why We Should Not Be Indifferent to Specification Choices for Difference-in-Differences," *Health Services Research* (50:4), pp. 1211–1235. (<https://doi.org/10.1111/1475-6773.12270>).
- Salgado-Montejo, A., Velasco, C., Olier, J. S., Alvarado, J., and Spence, C. 2014. "Love for Logos: Evaluating the Congruency between Brand Symbols and Typefaces and Their Relation to Emotional Words," *Journal of Brand Management* (21:7), pp. 635–649. (<https://doi.org/10.1057/bm.2014.29>).
- Segev, E. 2020. "Crowdsourcing Contests," *European Journal of Operational Research* (281:2), pp. 241–255. (<https://doi.org/10.1016/j.ejor.2019.02.057>).
- Seidel, S., Recker, J., and Brocke, J. vom. 2013. "Sensemaking and Sustainable Practicing: Functional Affordances of Information Systems in Green Transformations," *Management Information Systems Quarterly* (37:4), pp. 1275–1299.
- Senney, G. T. 2019. "The Geography of Lender Behavior in Peer-to-Peer Loan Markets," *The Journal of Industrial Economics* (67:1), pp. 1–36. (<https://doi.org/10.1111/joie.12195>).
- Shao, B., Shi, L., Xu, B., and Liu, L. 2012. "Factors Affecting Participation of Solvers in Crowdsourcing: An Empirical Study from China," *Electronic Markets* (22:2), pp. 73–82. (<https://doi.org/10.1007/s12525-012-0093-3>).
- Shortstack.com. 2018 How do I Run a Twitter Contest? Retrieved May 30, 2019, <https://www.shortstack.com/blog/how-to-run-a-twitter-contest/>
- Smith, A. N., Fischer, E., and Yongjian, C. 2012. "How Does Brand-Related User-Generated Content Differ across YouTube, Facebook, and Twitter?," *Journal of Interactive Marketing* (26:2), pp. 102–113. (<https://doi.org/10.1016/j.intmar.2012.01.002>).

- Smith, J. A., and Todd, P. E. 2005. "Does Matching Overcome LaLonde's Critique of Nonexperimental Estimators?," *Journal of Econometrics* (125:1–2), Elsevier, pp. 305–353.
- Socialmediatoday.com. 2016. Twitter's Algorithm Timeline Switches to Opt-Out, Coming to All Users. Retrieved May 30, 2019, <https://www.socialmediatoday.com/social-networks/twitters-algorithm-timeline-switches-opt-out-coming-all-users>
- Sproutsocial.com. 2016. How to Run an Actual Engaging Twitter Contest. Retrieved May 10, 2019, <https://sproutsocial.com/insights/twitter-contest/>
- Sproutsocial.com. 2018. How to Use Hashtag Marketing to Dramatically Boost Brand Awareness. Retrieved September 14, 2018, <https://sproutsocial.com/insights/hashtag-marketing/>
- Strong, D., Volkoff, O., Johnson, S., Pelletier, L., Tulu, B., Bar-On, I., Trudel, J., and Garber, L. 2014. "A Theory of Organization-EHR Affordance Actualization," *Journal of the Association for Information Systems* (15:2). (<https://aisel.aisnet.org/jais/vol15/iss2/2>).
- Sul, H., Dennis, A. R., and Yuan, L. I. 2014. "Trading on Twitter: The Financial Information Content of Emotion in Social Media," in *2014 47th Hawaii International Conference on System Sciences*, , January, pp. 806–815. (<https://doi.org/10.1109/HICSS.2014.107>).
- Sundstrom, B., and Levenshus, A. B. 2017. "The Art of Engagement: Dialogic Strategies on Twitter," *Journal of Communication Management* (21:1), pp. 17–33. (<https://doi.org/10.1108/JCOM-07-2015-0057>).
- Susarla, A., Oh, J.-H., and Tan, Y. 2011. "Social Networks and the Diffusion of User-Generated Content: Evidence from YouTube," *Information Systems Research* (23:1), pp. 23–41. (<https://doi.org/10.1287/isre.1100.0339>).
- Swani, K., Milne, G. R., Cromer, C., and Brown, B. P. 2013. "Fortune 500 Companies' Use of Twitter Communications: A Comparison between Product and Service Tweets," *International Journal of Integrated Marketing Communications* (5:2), pp. 47–56.
- Tilson, D., Lyytinen, K., and Sørensen, C. 2010. "Research Commentary—Digital Infrastructures: The Missing IS Research Agenda," *Information Systems Research* (21:4), pp. 748–759. (<https://doi.org/10.1287/isre.1100.0318>).
- Treem, J. W., and Leonardi, P. M. 2012. "Social Media Use in Organizations: Exploring the Affordances of Visibility, Editability, Persistence, and Association," SSRN Scholarly Paper No. ID 2129853, SSRN Scholarly Paper,

Rochester, NY: Social Science Research Network.
(<https://papers.ssrn.com/abstract=2129853>).

- Van den Bergh, J. C. J. M. 2008. "Optimal Diversity: Increasing Returns versus Recombinant Innovation," *Journal of Economic Behavior & Organization* (68:3), pp. 565–580. (<https://doi.org/10.1016/j.jebo.2008.09.003>).
- Volkoff, O., and Strong, D. M. 2013. "Critical Realism and Affordances: Theorizing IT-Associated Organizational Change Processes," *MIS Q.* (37:3), pp. 819–834. (<https://doi.org/10.25300/MISQ/2013/37.3.07>).
- Voorveld, H. A. M., Araujo, T., Bernritter, S. F., Rietberg, E., and Vliegthart, R. 2018. "How Advertising in Offline Media Drives Reach of and Engagement with Brands on Facebook," *International Journal of Advertising* (37:5), pp. 785–805. (<https://doi.org/10.1080/02650487.2018.1454703>).
- Wang, Q., Li, B., and Singh, P. V. 2018. "Copycats vs. Original Mobile Apps: A Machine Learning Copycat-Detection Method and Empirical Analysis," *Information Systems Research* (29:2), INFORMS, pp. 273–291. (<https://doi.org/10.1287/isre.2017.0735>).
- Wang, W., and Siau, K. 2019. "Artificial Intelligence, Machine Learning, Automation, Robotics, Future of Work and Future of Humanity: A Review and Research Agenda," *Journal of Database Management (JDM)* (30:1), IGI Global, pp. 61–79.
- Wang, Z., and Bovik, A. C. 2009. "Mean Squared Error: Love It or Leave It? A New Look at Signal Fidelity Measures," *IEEE Signal Processing Magazine* (26:1), pp. 98–117. (<https://doi.org/10.1109/MSP.2008.930649>).
- Webb, M. 2019. "The Impact of Artificial Intelligence on the Labor Market," SSRN Scholarly Paper No. ID 3482150, SSRN Scholarly Paper, Rochester, NY: Social Science Research Network, November 6. (<https://doi.org/10.2139/ssrn.3482150>).
- Wen, Z., and Lin, L. 2016. "Optimal Fee Structures of Crowdsourcing Platforms," *Decision Sciences* (47:5), pp. 820–850. (<https://doi.org/10.1111/deci.12201>).
- Wikihow.com. 2019. How to Promote with Twitter Hashtags. Retrieved May 30, 2019, <https://www.wikihow.com/Promote-with-Twitter-Hashtags>
- Windmeijer, F. 2005. "A Finite Sample Correction for the Variance of Linear Efficient Two-Step GMM Estimators," *Journal of Econometrics* (126:1), pp. 25–51.
- Wolbring, G., and Yumakulov, S. 2014. "Social Robots: Views of Staff of a Disability Service Organization," *International Journal of Social Robotics* (6:3), pp. 457–468. (<https://doi.org/10.1007/s12369-014-0229-z>).

- Wooten, J. O., and Ulrich, K. T. 2017. "Idea Generation and the Role of Feedback: Evidence from Field Experiments with Innovation Tournaments," *Production and Operations Management* (26:1), pp. 80–99. (<https://doi.org/10.1111/poms.12613>).
- Wooten, J., and Ulrich, K. 2015. "The Impact of Visibility in Innovation Tournaments: Evidence From Field Experiments," *Management Papers*. (<https://doi.org/10.2139/ssrn.2214952>).
- Wu, C., Yu, X., and Zheng, Y. 2020. "The Spillover Effect of Financial Information in Mergers and Acquisitions," *The British Accounting Review*, p. 100879. (<https://doi.org/10.1016/j.bar.2019.100879>).
- Wu, Z., and Salomon, R. 2016. "Does Imitation Reduce the Liability of Foreignness? Linking Distance, Isomorphism, and Performance," *Strategic Management Journal* (37:12), pp. 2441–2462. (<https://doi.org/10.1002/smj.2462>).
- Yang, J., She, D., and Sun, M. 2017. "Joint Image Emotion Classification and Distribution Learning via Deep Convolutional Neural Network," in *IJCAI*, pp. 3266–3272.
- Yang, Y., Chen, P., and Banker, R. D. 2010. Impact of Past Performance and Strategic Bidding on Winner Determination of Open Innovation Contest. In *Workshop on Information Systems and Economics (WISE 2010)* (Vol. 2010, pp. 11-2)
- Yoo, Y., Boland, R. J., Lyytinen, K., and Majchrzak, A. 2012. "Organizing for Innovation in the Digitized World," *Organization Science* (23:5), pp. 1398–1408. (<https://doi.org/10.1287/orsc.1120.0771>).
- You, Q., Luo, J., Jin, H., and Yang, J. 2016. "Building a Large Scale Dataset for Image Emotion Recognition: The Fine Print and the Benchmark," in *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence, AAAI'16*, Phoenix, Arizona: AAAI Press, pp. 308–314. (<http://dl.acm.org/citation.cfm?id=3015812.3015857>).
- Youn, H., Bettencourt, L. M. A., Strumsky, D., and Lobo, J. 2015. "Invention as a Combinatorial Process: Evidence from U.S. Patents," *Journal of The Royal Society Interface* (12:106), p. 20150272. (<https://doi.org/10.1098/rsif.2015.0272>).
- Yu, H., and Winkler, S. 2013. "Image Complexity and Spatial Information," in *2013 Fifth International Workshop on Quality of Multimedia Experience (QoMEX)*, IEEE, pp. 12–17.
- Zeng, X., and Wei, L. 2012. "Social Ties and User Content Generation: Evidence from Flickr," *Information Systems Research* (24:1), pp. 71–87. (<https://doi.org/10.1287/isre.1120.0464>).

- Zhang, J., Yan, Y., and Guan, J. 2019. "Recombinant Distance, Network Governance and Recombinant Innovation," *Technological Forecasting and Social Change* (143), pp. 260–272. (<https://doi.org/10.1016/j.techfore.2019.01.022>).
- Zhang, S., Singh, P. V., and Ghose, A. 2019. "A Structural Analysis of the Role of Superstars in Crowdsourcing Contests," *Information Systems Research* (30:1), INFORMS, pp. 15–33.
- Zheng Zhou, K. 2006. "Innovation, Imitation, and New Product Performance: The Case of China," *Industrial Marketing Management* (35:3), New Product Development in Asia, pp. 394–402. (<https://doi.org/10.1016/j.indmarman.2005.10.006>).
- Zhou Wang, Bovik, A. C., Sheikh, H. R., and Simoncelli, E. P. 2004. "Image Quality Assessment: From Error Visibility to Structural Similarity," *IEEE Transactions on Image Processing* (13:4), pp. 600–612. (<https://doi.org/10.1109/TIP.2003.819861>).
- Zittrain, J. 2006. "The Generative Internet," *Harvard Law Review*. (<https://doi.org/10.1145/1435417.1435426>).