Università Commerciale "Luigi Bocconi"

PhD School

PhD program in: Business Administration and Management

Cycle: XXXVI

Disciplinary Field (code): SECS - P/08

Strategic Insight from Unstructured Data: Essays on Innovation, Brand Communities, and Entrepreneurship

Advisor: Gaia Rubera

PhD Thesis by

Margherita Caprara

ID number: 3039889

Year: 2026

Abstract

This dissertation explores how unstructured data, textual and visual content shared by individuals and organizations, can offer valuable insights into strategic priorities, brand communities, and entrepreneurial activity. Leveraging recent advances in artificial intelligence, particularly natural language processing and computer vision, the three essays apply and extend scalable approaches to analyze patent texts, social media posts, and image-based signals. Together, the studies illustrate how multi-modal unstructured data can uncover overlooked patterns in innovation, brand positioning, and entrepreneurship.

The first chapter examines whether and when female leadership promotes innovation that targets women's needs. Using a text-based measure of female-focused innovation constructed from pharmaceutical patent titles and abstracts, the study finds that female representation in top management teams is not sufficient: only under specific personal conditions, such as parenting daughters, do leaders redirect strategic attention toward more inclusive innovation.

The second chapter proposes a new framework for analyzing brand communities on Instagram. Combining computer vision and community detection, it identifies lifestyle-based consumer segments based on how logos and hashtags co-occur in millions of posts. The analysis of Adidas and Nike across global cities reveals both shared and localized brand meanings, offering a methodological contribution to brand positioning and community analysis.

The third chapter connects narcissistic self-presentation on Instagram to entrepreneurial

activity across U.S. cities. Using a deep learning model to identify selfies as a proxy for non-pathological narcissism, the study finds that selfie prevalence is positively associated with startup quantity but not startup quality. The findings contribute to regional entrepreneurship and psychological economics by introducing a scalable, image-based behavioral signal of entrepreneurial orientation.

Across these chapters, the dissertation advances the marketing field by showing how multi-modal, unstructured data, encompassing patent texts, social media images, and hashtags, can be utilized to investigate innovation focus, brand communities, and the drivers of entrepreneurship.

Contents

Introduction **7**

1 The Gender-Innovation Paradox: Female Leadership, Parenting, and Female-Focused Innovation **15**

1.1 Introduction 17

1.2 Theoretical Background 19

1.2.1 Top Management Team and Innovation Focus 19

1.2.2 Women in leadership positions 23

1.2.3 The role of environment: The parenting effect 25

1.3 Hypotheses 27

1.3.1 The relationship between female influence in the TMT and female-focused innovation 28

1.3.2 The relationship between parenting and female-focused innovation . 30

1.3.3 The moderating effect of parenting 33

1.4 Method 35

1.4.1 Sample 35

1.4.2 Measures 36

1.4.3 Estimation strategy 39

1.5 Results 45

1.5.1 Main analyses 45

1.5.2 Robustness analyses . . . 50
1.6 Discussion . . . 51
1.6.1 Theoretical contributions . . . 51
1.6.2 Managerial implications . . . 54
1.6.3 Limitations and Directions for Future Research . . . 56
Appendices . . . 57
A1 Appendix: Construction of the Final Dataset . . . 57
A2 Appendix: Face validity of the measure of female-focused innovation . . . 58

2 A Day in the Life: Mapping Brand DNA and Local Identities Through Multimodal User-Generated Content 71

2.1 Introduction . . . 73
2.2 Literature on User Generated Content . . . 76
2.3 Data and Methods . . . 80
2.3.1 Instagram Images Dataset . . . 80
2.3.2 Brand Logo Detection . . . 81
2.3.3 Community Detection Algorithm . . . 83
2.3.4 Brand Community Labeling and Description . . . 85
2.4 Application of MBCA . . . 87
2.4.1 Within-city and Between-brands . . . 88
2.4.2 Within-brand and Between-cities . . . 90
2.4.3 Brand Logo versus Brand Mention . . . 93
2.5 Conclusion . . . 95
2.5.1 Theoretical Implications . . . 95
2.5.2 Managerial Implications . . . 98
2.5.3 Limitations . . . 103
Appendices . . . 105

A1 Appendix: Between-city and Within-brand . . . 105
A1.1 Nike Communities in San Francisco and Los Angeles . . . 105
A1.2 Nike Communities Across Cities . . . 107
A2 Appendix: Prompt Documentation . . . 108
A3 Appendix: Community Campaigns . . . 109

3 Assessing Entrepreneurial Activity through Deep Learning: A Geographic Analysis Using Instagram Data 115
3.1 Introduction . . . 117
3.2 Theoretical Background . . . 118
3.3 Data and Method . . . 120
3.3.1 Selfie Detection . . . 121
3.3.2 Entrepreneurial Quantity . . . 121
3.3.3 Entrepreneurial Quality . . . 122
3.3.4 Control Variables . . . 123
3.3.5 Estimation Strategy . . . 124
3.4 Results . . . 125
3.4.1 Model 1: Entrepreneurial Quantity . . . 125
3.4.2 Model 2: Entrepreneurial Quality . . . 126
3.5 Discussion and Conclusion . . . 126
Appendix . . . 129
A1 Appendix: Survey Validation of the Narcissism Proxy . . . 129

Acknowledgements

First and foremost, I would like to express my deepest gratitude to my advisor, Gaia Rubera. I first met her during my master's, when I enrolled in her course that combined machine learning and marketing — a decision that ultimately shaped my academic path. Her guidance, high standards, and belief in meaningful research have been instrumental throughout this journey, and I am truly grateful for her mentorship.

I am also thankful to Paola Cillo for her insight and example, to Sara Valentini for her support and for giving me my first opportunity to teach, and to Francesco Grossetti for helping me build the technical skills I needed early on.

I am grateful to the Bocconi marketing department for providing a rigorous and supportive environment throughout these years.

I extend my thanks to my fellow PhD students, both those with whom I shared daily life and those I met at conferences. Their support and camaraderie made this journey far richer.

To my parents, my sister Benedetta, and my brothers Pietro and Carlo — thank you for your constant support, trust, and encouragement. And to Alberto — for the patience, care, and presence through every phase of this journey — I am deeply grateful.

This PhD is only the beginning. But as my favorite TV show says: *Nothing in this world worth having comes easy.*

Introduction

In an increasingly digital and visual world, marketers and scholars alike are confronted with a fundamental shift in the availability and nature of data. Beyond traditional structured data sets, consumers, companies, and entrepreneurs are continuously generating vast streams of unstructured content, including text, images, hashtags, and interactions, that encode attitudes, priorities, and decisions in subtle, often indirect ways. At the same time, methodological advancements in artificial intelligence, particularly in deep learning and natural language processing, have unlocked new possibilities for extracting insights from this content at scale.

This dissertation contributes to the growing body of marketing scholarship that leverages unstructured data to examine economic and social phenomena. Across three empirical studies, it shows how AI-powered methods can uncover latent strategic, cultural, and psychological dynamics by analyzing what firms patent and how visual and textual signals co-occur in user-generated content. The chapters are united by a common methodological approach: using text analysis, computer vision, and community detection. They also share a theoretical ambition: to illuminate the drivers of innovation focus, brand communities, and entrepreneurial activity by analyzing the visual and textual content that individuals and organizations deliberately create and share.

Chapter 1

The first chapter, *The Gender-Innovation Paradox: Female Leadership, Parenting, and Female-Focused Innovation*, examines whether and under what conditions leadership diversity is associated with more inclusive innovation outcomes. Drawing on upper echelons theory (Hambrick and Mason, 1984) and the attention-based view of the firm (Ocasio, 1997), the study shows that the gender composition of the top management team (TMT), the relative distribution of power within it, and the personal life experiences of its members, particularly parenting, are jointly related to the degree to which firms engage in female-focused innovation.

Using a panel of pharmaceutical firms from 2000 to 2019, the study employs a text-based measure of female-focused innovation extracted from patent titles and abstracts. The results challenge the prevailing assumption that increased female representation in leadership naturally leads to greater attention to women's needs. In fact, greater female influence in the TMT is, on average, associated with lower levels of female-focused innovation. However, this relationship is moderated by parenting: it is attenuated or reversed when more TMT members have daughters, and becomes more negative when more have sons.

These findings contribute to research on gender and innovation by showing that leadership diversity alone may not be sufficient to redirect innovation toward underserved groups. By integrating structural and experiential dimensions of executive life, the study demonstrates how personal experiences, particularly parenting daughters, can make gender-related concerns more prominent in strategic decision-making. In doing so, it extends existing work on gender and innovation focus (Koning et al., 2021; Einiö et al., 2019) and offers a more nuanced view of how top executives allocate attention in shaping inclusive innovation agendas.

This chapter lays the groundwork for the dissertation's broader methodological ap-

proach by demonstrating how unstructured textual data can be transformed into strategic insight. Patent documents, often treated as technical or administrative artifacts, are employed as expressions of firms' innovation priorities and market focus. By applying natural language processing techniques to this content and linking it to demographic data on executives, the study shows how text-based analysis can reveal systematic patterns in how firms attend to underserved market segments.

Chapter 2

The second chapter, *A Day in the Life: Mapping Brand DNA and Local Identities Through Multimodal User-Generated Content*, develops a method for analyzing brand communities through deep learning-based computer vision and network analysis. The study identifies brand logos in Instagram posts and clusters the associated hashtags into communities, offering marketers a scalable way to uncover lifestyle-based consumer segments. Applied to Adidas and Nike across eight cities, the method reveals both global consistencies and local adaptations in how consumers relate to these brands.

A key contribution is the identification of "context bias," wherein consumers only mention brands explicitly in posts that align with expected use contexts (e.g., fitness-focused communities for athletic brands). Relying solely on brand mentions thus overlooks broader patterns of brand-related behavior. This paper proposes that by visually detecting brand logos, even in the absence of brand hashtags, and analyzing surrounding hashtags, marketers can recover a richer and more nuanced representation of the brand's audience. The chapter contributes to research on user-generated content and automated market research by introducing a scalable method to analyze brand communities through visual and textual data (Lee and Bradlow, 2011; Netzer et al., 2012; Rust et al., 2021). By identifying both global themes and local adaptations, it offers marketers a tool to understand how brand associations vary across cities and contexts, enabling more targeted

and context-aware strategies.

Moreover, by revealing behavioral patterns such as context bias, where consumers selectively mention brands in use-aligned posts, the method expands the scope of visual UGC analysis (Dzyabura and Peres, 2021; Klostermann et al., 2018; Liu et al., 2020) and provides a richer and more multidimensional view of brand engagement.

Chapter 3

The final chapter, *Assessing Entrepreneurial Activity through Deep Learning: A Geographic Analysis Using Instagram Data*, bridges collective psychology, regional entrepreneurship, and computer vision. It investigates whether narcissistic self-presentation, proxied by the prevalence of selfies, predicts startup formation across US cities. Narcissism, a stable personality trait characterized by self-importance and a need for admiration (Campbell et al., 2004; Engelen et al., 2016), has been linked to entrepreneurial orientation and leadership style. However, most studies rely on self-reported data or elite samples (e.g., CEOs), limiting their generalizability.

This paper introduces a scalable behavioral proxy: the proportion of Instagram images classified as selfies using YOLO v3, a state-of-the-art deep learning model (Redmon and Farhadi, 2018). Merged with Crunchbase data on startups and the Entrepreneurial Quality Index (EQI) (Andrews et al., 2022), the analysis reveals that selfie prevalence is positively associated with the number of startups founded, but not with startup quality. These findings suggest that narcissistic expression may catalyze action (Grijalva et al., 2015; Leung et al., 2021), but does not guarantee performance. Methodologically, the chapter contributes to the growing use of big data in regional psychology (Obschonka et al., 2020) and highlights how computer vision can be applied to sociocultural phenomena.

This dissertation contributes to research on innovation strategy, user-generated con-

tent, and entrepreneurship by showing how multi-modal unstructured data, spanning patent texts, social media images, and hashtags, can be leveraged to examine organizational priorities and market-level patterns. The first essay extends the work on gender and innovation by showing that female leadership is not automatically associated with female-focused innovation, and that parenting experiences, particularly having daughters, play a moderating role (Koning et al., 2021; Einiö et al., 2019). The second essay introduces a method for analyzing brand communities using computer vision and community detection, revealing how brand meanings vary globally and locally in Instagram content (Dzyabura and Peres, 2021; Klostermann et al., 2018; Liu et al., 2020). The third essay links narcissistic self-presentation, proxied by selfies, to startup activity across U.S. cities, contributing to research on regional entrepreneurial orientation (Grijalva et al., 2015; Obschonka et al., 2020).

Methodologically, the dissertation develops approaches to analyze visual and textual data at scale using tools from computer vision, natural language processing, and econometrics. These approaches make it possible to study strategic priorities, brand associations, and collective personality traits using data that standard marketing methods typically overlook.

References

Andrews, R. J., Fazio, C., Guzman, J., Liu, Y., and Stern, S. (2022). Reprint of "the startup cartography project: measuring and mapping entrepreneurial ecosystems". *Research Policy*, 51(9):104581.

Campbell, W. K., Goodie, A. S., and Foster, J. D. (2004). Narcissism, confidence, and risk attitude. *Journal of behavioral decision making*, 17(4):297–311.

Dzyabura, D. and Peres, R. (2021). Visual elicitation of brand perception. *Journal of Marketing*, 85(4):44–66.

Einiö, E., Feng, J., and Jaravel, X. (2019). Social push and the direction of innovation. *Available at SSRN 3383703*.

Engelen, A., Neumann, C., and Schmidt, S. (2016). Should entrepreneurially oriented firms have narcissistic ceos? *Journal of Management*, 42(3):698–721.

Grijalva, E., Harms, P. D., Newman, D. A., Gaddis, B. H., and Fraley, R. C. (2015). Narcissism and leadership: A meta-analytic review of linear and nonlinear relationships. *Personnel Psychology*, 68(1):1–47.

Hambrick, D. C. and Mason, P. A. (1984). Upper echelons: The organization as a reflection of its top managers. *Academy of management review*, 9(2):193–206.

Klostermann, J., Plumeyer, A., Böger, D., and Decker, R. (2018). Extracting brand information from social networks: Integrating image, text, and social tagging data. *International Journal of Research in Marketing*, 35(4):538–556.

Koning, R., Samila, S., and Ferguson, J.-P. (2021). Who do we invent for? Patents by women focus more on women's health, but few women get to invent. *Science*, 372(6548):1345–1348.

Lee, T. Y. and Bradlow, E. T. (2011). Automated marketing research using online customer reviews. *Journal of Marketing Research*, 48(5):881–894.

Leung, Y. K., Franken, I., Thurik, R., Driessen, M., Kamei, K., Torres, O., and Verheul, I. (2021). Narcissism and entrepreneurship: Evidence from six datasets. *Journal of Business Venturing Insights*, 15:e00216.

Liu, L., Dzyabura, D., and Mizik, N. (2020). Visual listening in: Extracting brand image portrayed on social media. *Marketing Science*, 39(4):669–686.

Netzer, O., Feldman, R., Goldenberg, J., and Fresko, M. (2012). Mine your own business: Market-structure surveillance through text mining. *Marketing Science*, 31(3):521–543.

Obschonka, M., Lee, N., Rodríguez-Pose, A., Eichstaedt, J. C., and Ebert, T. (2020). Big data methods, social media, and the psychology of entrepreneurial regions: capturing cross-county personality traits and their impact on entrepreneurship in the usa. *Small Business Economics*, 55:567–588.

Ocasio, W. (1997). Towards an attention-based view of the firm. *Strategic Management Journal*, 18(S1):187–206.

Redmon, J. and Farhadi, A. (2018). Yolov3: An incremental improvement. *arXiv preprint arXiv:1804.02767*.

Rust, R. T., Rand, W., Huang, M.-H., Stephen, A. T., Brooks, G., and Chabuk, T. (2021). Real-time brand reputation tracking using social media. *Journal of Marketing*, 85(4):21–43.

Chapter 1

The Gender-Innovation Paradox: Female Leadership, Parenting, and Female-Focused Innovation

Margherita Caprara[1], **Paola Cillo**[2], **Gaia Rubera**[1]

[1]Department of Marketing, Bocconi University

[2]Department of Management and Technology, Bocconi University

Female consumers control over $31.8 trillion in annual global spending, yet they remain systematically underserved across industries. While common wisdom suggests that increasing women's representation in leadership should foster more female-focused innovation, we challenge this assumption. Drawing on the attention-based view of the firm and upper echelons theory, we argue that identity alone is not enough, but top executives' attention is also shaped by context and personal experience. Analyzing pharmaceutical firms from 2000 to 2019, we find that greater female influence in the top management team (TMT) is, on average, associated with lower levels of female-focused innovation. This counterintuitive effect is conditional: it is weakened, and even reversed, when more TMT members have daughters. In contrast, parenting sons exacerbates the negative association. These findings offer a novel behavioral lens on managerial attention, revealing how emotionally salient life experiences shape firms' responsiveness to underserved markets. We contribute to marketing scholarship by showing when leadership diversity translates into inclusive innovation. Ultimately, unlocking the full potential of the female consumer market requires not just more female influence in the TMT, but organizational contexts that heighten male leaders' attentional sensitivity to gender-related concerns.

Keywords: innovation; underserved market; attention-based view; executive decisions

1.1 Introduction

A growing body of research in marketing has explored how brands can serve underrepresented consumers (D'Angelo et al., 2024). Yet, many of these initiatives appear more symbolic than strategic, aiming to signal inclusivity rather than uncovering meaningful opportunities for growth. Not surprisingly, in response to political pushback, several brands have begun rolling back diversity, equity, and inclusion efforts (Sands and Ferraro, 2025). However, in an increasingly competitive global economy, identifying and serving underserved markets offers brands a powerful avenue for growth and differentiation.

Among these underserved segments, female consumers stand out. Despite controlling an estimated $31.8 trillion in annual global spending[1], women remain systematically overlooked across innovation domains. In healthcare, for instance, only 5% of annual R&D funding targets female-specific conditions (Kemble et al., 2022), and gynecologic cancers receive far less support relative to their lethality than diseases like prostate cancer (Spencer et al., 2019). Gender gaps persist in several industries, such as automotive (Bose et al., 2011), consumer electronics (Cifor and Garcia, 2019), and space suits (Mickey and Smith-Doerr, 2022). These disparities represent not only social inequities but also untapped market potential.

Building on the notion that top executives' attention shapes innovation (Yadav et al., 2007), recent work suggests that the gender gap in commercialized products may reflect a gender gap in decision-making influence (Koning et al., 2021). Accordingly, both academic and policy efforts have promoted gender diversity in leadership as a corrective to historical bias in innovation (Graves Jr et al., 2022; Iyer, 2020; Mickey and Smith-Doerr, 2022). Interventions such as executive quotas, board mandates, and targeted recruitment are premised on the idea that women in leadership will direct greater attention to underserved female consumers. However, existing research has focused primarily on the

[1] https://shorturl.at/UyDgJ (accessed April 29th, 2025.)

volume of innovation (You et al., 2020), rather than its focus. Whether female leadership actually associates with higher innovation directed at women's needs (i.e., *female-focused innovation*) remains an open question for scholars and policy makers alike. Addressing this gap is essential for understanding if and when female representation helps firms cater to the needs of unserved customers and for developing more effective initiatives to increase innovation inclusivity.

To address this question, we bring together upper echelons theory (Hambrick and Mason, 1984) and the attention-based view of the firm (Ocasio, 1997) to examine how the composition of the Top Management Team (TMT) associates with female-focused innovation. We propose that attention to female customers can arise from two primary sources. First, identity-based mechanisms, such as the homophily effect, suggest that women in leadership are more likely to prioritize issues that reflect their own lived experiences. Yet, this influence is not guaranteed: as the queen bee hypothesis highlights, women in male-dominated environments may suppress advocacy for female concerns to maintain their legitimacy and professional standing (Derks et al., 2015, 2016). Second, drawing on the attention-based view (Ocasio, 1997), we argue that exposure to women in one's social environment—especially through parenting daughters—can make gender-related needs more cognitively salient even for male executives. In contrast, executives with sons may remain anchored in more traditional, male-coded innovation domains.

Using these theoretical foundations, we develop a novel framework that links female influence in the TMT (Srivastava et al., 2023), the gender of executives' children, and female-focused innovation. We test both direct and interaction effects. This study yields three main findings. First, female influence in the TMT is negatively associated with female-focused innovation—contrary to prevailing assumptions that women in leadership naturally prioritize female concerns. Second, a higher proportion of TMT members with daughters is associated with greater female-focused innovation, whereas a higher proportion with sons is associated with less. Third, we observe that parenting experi-

ences systematically moderate the relationship between female influence in the TMT and female-focused innovation: as the proportion of TMT members with daughters increases, the negative association is attenuated and eventually becomes positive; conversely, as the proportion with sons increases, the negative association grows stronger. Collectively, these results indicate that greater female influence in the TMT is not consistently associated with more female-focused innovation; rather, it is contingent on the parenting experiences of TMT members.

This study offers three key theoretical contributions. First, we challenge the prevailing assumption that greater female influence in top leadership naturally drives female-focused innovation (Graves Jr et al., 2022; Mickey and Smith-Doerr, 2022). Our findings indicate that this relationship appears contingent on the broader social context (both at work and at home), calling into question the presumed direct link between gender diversity and female-focused innovation. Second, we extend the attention-based view of the firm by empirically demonstrating that executives' attention is not only situated in organizational structures, but also socially and emotionally shaped by relevant personal experiences such as parenting daughters or sons, which are associated with distinct patterns of female-focused innovation. Third, we extend the gender-innovation literature (Koning et al., 2020; Einiö et al., 2019) by shifting the focus upstream—from inventors to firm leadership—and show that characteristics of TMT members are associated with how firms allocate attention to female-focused innovation, highlighting the strategic role of TMT in shaping inclusive innovation agendas.

1.2 Theoretical Background

1.2.1 Top Management Team and Innovation Focus

A top management team (TMT) is the small group of executives responsible for charting firm strategy (Srivastava et al., 2023). It plays a central role in setting a firm's in-

novation strategy by prioritizing innovation areas, allocating resources accordingly, and defining key customer segments (Harmancioglu et al., 2010; Talke et al., 2010). These decisions depend heavily on TMT members' awareness of external opportunities and their interpretation of environmental signals (Yadav et al., 2007). This awareness emerges through attention, which filters and frames external information (Li et al., 2013). Consequently, the type of innovations a firm pursues is largely a reflection of the attentional focus of its TMT members.

To understand how female-focused innovation emerges within firms, we draw on two complementary theories: upper echelons theory and the attention-based view of the firm. The former emphasizes how executives' identities shape strategic outcomes, while the latter explains how attention is selectively and situationally directed.

Upper echelons theory and attention-based view

Upper echelons theory posits that a firm's outcomes largely reflect its senior executives' interpretation of the environment and the amount of attention they pay to it (Hambrick and Mason, 1984). It maintains that executives' identities—reflected in both demographic attributes (e.g., gender, age) and experiential traits (e.g., functional background, professional expertise)—serve as valid proxies for how they interpret their environment and make strategic choices. For example, board members with marketing experience tend to focus on customer concerns, buffering the negative effects of mergers on customer satisfaction Umashankar et al. (2022). Similarly, marketing executives are more attuned to consumer needs and trends, thereby fostering customer-centric innovation (Brower and Nath, 2018). Crucially, this theory contends that firm outcomes are better explained by the characteristics of the entire TMT rather than those of the CEO alone (Hambrick, 2007).

The attention-based view complements upper echelons theory by examining how identity and context jointly determine what executives notice and prioritize (Ocasio, 1997,

2011). Because attention is inherently limited, executives must filter and prioritize what they notice, interpret, and act upon. This filtering process means that what executives pay attention to plays a central role in shaping a firm's resource allocation and its strategic direction (Umashankar et al., 2022). For instance, executives who pay more attention to innovation allocate more resources toward innovation-related activities (Yadav et al., 2007). Thus, rather than assuming that executives respond uniformly to the environment, this view highlights that they attend selectively to certain issues over others (Day and Nedungadi, 1994; Harmancioglu et al., 2010; Li et al., 2013; Yadav et al., 2007).

To explain how this stimulus filtering occurs, the attention-based view distinguishes between individual-level attention and *situated* attention. The former is rooted in a decision maker's personal identity. The latter–which Ocasio (1997, p. 204) considers the "central contribution" of his theory–emphasizes that attention largely depends on the context in which an individual is embedded at a given moment.

In this view, situational factors influence which issues are noticed, framed as legitimate, and acted upon. These factors may be internal to the organization or part of the broader social environment—what Ocasio (1997) refers to as "environmental embeddedness." Thus, attention does not depend solely on the identity of the executives but also on whom they interact with.

Taken together, upper echelons theory and the attention-based view suggest that TMT members are more likely to focus on issues that are personally relevant, which influences firm outcomes, including innovation (Li et al., 2013; Ocasio, 2011). As a result, firms tend to innovate in domains that reflect the attentional schemas and values of their top executives. We argue that the extent to which firms engage in female-focused innovation depends on how salient women-related issues are for the TMT members. Drawing on upper echelons theory, we propose that female-focused innovation may be associated with the identity of its members; specifically, the degree of *female influence within the TMT* (Srivastava et al., 2023). In line with the attention-based view's emphasis on attention as

a situated and selective process, we further propose that contextual cues—particularly the *gender of TMT members' children*—may be associated with female-focused innovation.

We now turn to define the construct of female influence within the TMT, which serves as a central element of our theoretical framework.

Female influence in the TMT

According to the attention-based view, attention is socially constructed through ongoing negotiation among actors with differing priorities (Ocasio, 2011).

Power plays a critical role in these processes (Brielmaier and Friesl, 2023). Within TMTs, power is not evenly distributed. Even in demographically diverse teams, influence is often concentrated within dominant coalitions (Cyert and March, 1963). As Finkelstein (1992) outlines, power members of the TMT may remain peripheral to core strategic conversations, even if numerically present. Consequently, the ability of women in the TMT to shape organizational attention and innovation outcomes depends more on their influence than on their numbers.

A related stream of work in corporate governance underscores the importance of critical mass. Torchia et al. (2011) find that the influence of women on boards becomes visible only when a minimum threshold is reached. Below this level, women are often perceived as tokens, face marginalization, and struggle to influence group dynamics. Once critical mass is reached, however, the group is no longer seen as a token presence but as a legitimate minority capable of altering strategic priorities—such as increasing firm innovation.

Building on these insights, recent research in marketing highlights the importance of female influence in the TMT as a more accurate predictor of strategic outcomes than representation alone. Srivastava et al. (2023) define it as the "degree of influence female executives exert over decision making in the TMT" (197), reflecting both formal authority and informal influence. They show that this influence is positively associated with customer orientation and long-term financial performance. Importantly, the study cau-

tions against treating gender diversity as a "body count" metric; what matters is whether women have the power to channel the attention of the TMT.

To understand whether female influence in the TMT relates to the prioritization of women-related issues, we turn to review the literature examining how women in leadership positions influence the prioritization of women-related issues within firms.

1.2.2 Women in leadership positions

The existing literature offers two contrasting perspectives on how women in leadership roles influences the prioritization of women-related issues within organizations: the homophily hypothesis and the queen bee hypothesis.

Homophily hypothesis

The *homphily* hypothesis posits that individuals are more likely to engage in activities—such as innovation, research, or decision-making—that reflect the characteristics and experiences of their own social identity, including gender, race, and socioeconomic background (Einiö et al., 2019; Koning et al., 2021; Kozlowski et al., 2022; Nielsen et al., 2018). In innovation contexts, this implies that individuals tend to focus on problems and design solutions that resonate with their lived experiences. For instance, Josephine Cochrane's invention of the dishwasher in 1886 exemplifies how gendered experiences can inspire novel product development (Einiö et al., 2019).

Empirical evidence supports the homophily hypothesis. Koning et al. (2020) find that female-majority teams are significantly more likely to develop inventions targeting women (18.5%) and less likely to develop those for male-focused inventions (6%), with the reverse pattern observed among male-majority teams. This tendency explains why, in a world where inventors are predominantly men, aggregate innovation remains skewed toward male-centered needs and interests. Homophily thus contributes to gendered disparities in research attention and product development. But the effects of homophily extend beyond

invention. Minoritized scientists often publish in disciplines and explore research questions that resonate with their gendered and racialized social identities (Kozlowski et al., 2022). Similarly, in organizational settings, greater female representation in top leadership has been linked to increased support for other women (Gould et al., 2018; Guldiken et al., 2019), especially at junior levels (Gorman, 2005). These cascading effects suggest that homophily influences what gets invented, studied, and promoted.

Beyond organizations and innovation, homophily has been shown to influence a wide range of outcomes, including citation patterns in academic publishing (Zhou et al., 2024), evaluations in the film industry (Stroube and Waguespack, 2024), and participation rates in clinical trials (Gupta, 2022). Collectively, these studies highlight how leader gender shapes attention, decision-making, and ultimately, innovation focus.

Queen bee hypothesis

In contrast, the queen bee hypothesis suggests that some women in senior roles may consciously distance themselves from women-related issues, particularly in male-dominated environments. Originally introduced by Staines et al. (1974), this phenomenon describes how high-status women may attempt to assimilate into masculine organizational cultures by downplaying gender identity and resisting initiatives aimed at supporting other women (Derks et al., 2016; Ellemers et al., 2004).

Laboratory studies suggest that heightened gender salience and concerns about being devalued through association with other women may lead high-status women to refrain from promoting or selecting other female colleagues (Duguid, 2011; Duguid et al., 2012). Anecdotal accounts echo these findings. For instance, in the 1990s, a group of female scientists at MIT discreetly formed a committee to investigate gender inequities. Molecular biologist Nancy Hopkins later recalled that "like me, these women scientists did not want to be seen as feminists."[2] More recently, in 2021, computer scientist Katie

[2] https://shorturl.at/f3LCo

Bouman—who led the development of the algorithm used to capture the first image of a black hole—expressed similar reluctance: "At first, I didn't want to be seen as a woman in science, I just wanted to be recognized as a scientist."[3] These accounts highlight a persistent tension: in male-dominated fields, women may fear that advocating for gender equity could undermine the recognition of their scientific achievements. These behaviors reflect what Derks et al. (2015) call "self-group distancing"—a coping strategy that arises not from personality traits, but from efforts to navigate gendered organizational bias. In their review, Derks et al. (2016) argue that queen bee behavior is a response to identity threat and anticipation of discrimination in male-dominated workplaces. Importantly, this response is not unique to women; similar distancing behaviors have been documented among other marginalized groups facing workplace stereotypes and devaluation (Derks et al., 2015; Fordham, 2008).

1.2.3 The role of environment: The parenting effect

While identity is a critical driver of managerial attention, it is not the only one. According to attention-based view, leaders' social environments also shape what they notice, prioritize, and ultimately act upon (Ocasio, 1997). Unlike identity-based mechanisms such as homophily and queen bee dynamics, situated attention emphasizes the role of context, particularly how exposure to others can make certain issues more salient.

The *female socialization* literature suggests that individuals internalize the perspectives and concerns of those in their immediate social environment, even when they do not share those identities themselves (Cronqvist and Yu, 2017; Einiö et al., 2019). In particular, exposure to women can increase men's awareness of gender-related challenges and lead them to prioritize issues relevant to women. Several empirical studies support this hypothesis. For instance, Truffa and Wong (2024) find that when U.S. universities transitioned from male-only to coeducational institutions, incumbent (primarily male)

[3] https://shorturl.at/TSZue

faculty began shifting their research focus toward topics related to sex differences and female-related concerns. The authors attribute this finding to the fact that the changing composition of classroom generated new discussion topics and altered course content, which influenced faculty research interests. Similarly, Einiö et al. (2019) show that male entrepreneurs exposed to female peers are more likely to develop products targeting female consumers.

This effect is especially pronounced in close, emotionally salient relationships, such as father-daughter dynamics. A growing body of literature finds that having girls profoundly influence parental values and behaviors, making males with a daughter more sensitive to female instances than other males. In support of this perspective, Washington (2008) shows that U.S. Congressmen with daughters are significantly more likely to support liberal positions, especially on reproductive rights. Glynn and Sen (2015) report similar findings among federal judges: male judges with daughters issue more liberal rulings on gender-related cases. In corporate contexts, the "*daughter effect*" is increasingly recognized as a meaningful driver of gender-equity decisions. CEO with daughters are more likely to invest in corporate social responsibility (Cronqvist and Yu, 2017) and appoint women to the boards (Dasgupta et al., 2018). Venture capital partners with daughters are more prone to hire female partners (Calder-Wang and Gompers, 2021), and male startup founders with daughters tend to appoint more women to leadership positions (Wu et al., 2024). Across these studies, the underlying mechanism is clear: exposure to daughters increases attentional sensitivity to women's challenges and enhances the likelihood of acting to support gender equity.

By contrast, the effect of parenting sons remains less explored, but emerging evidence suggests that sons may have a very different—or even opposite—effect. Oswald and Powdthavee (2010) find that British parents with sons tend to become more politically conservative, while those with daughters become more liberal. Similarly, fathers of sons are more likely to endorse or engage in patriarchal behaviors (Kabátek and Ribar, 2021;

Raley and Bianchi, 2006), and are less motivated to engage with gender-equity concerns because they do not perceive them as a threat to their children (Warner and Steel, 1999).

Taken together, these findings suggest that the influence of parenting on managerial attention is gender-contingent: daughters increase sensitivity to women-related issues, while sons may reinforce traditional gender norms or lead to a focus on male-typed concerns. From the lens of the attention-based view, parenting daughters or sons provides situational cues that differentially shape what issues executives find salient, legitimate, and worth acting on.

1.3 Hypotheses

Our literature review highlights two key antecedents of TMT attention to women-related issues: (1) identity-based mechanisms, such as homophily and queen bee dynamics, and (2) contextual influences arising from executives' personal environments. In particular, the "daughter effect" exemplifies how situated attention can sensitize executive, especially men, to female-specific concerns. Conversely, parenting sons may reinforce traditional gender norms and maintain existing attentional patterns, thereby diminishing the salience of female-specific concerns in strategic decision-making. Building on this theoretical foundation, we develop a set of hypotheses exploring how (a) female influence within the TMT, (b) the gender of TMT members' children, and (c) their interaction associate with female-focused innovation. Figure 1.1 depicts our conceptual framework.

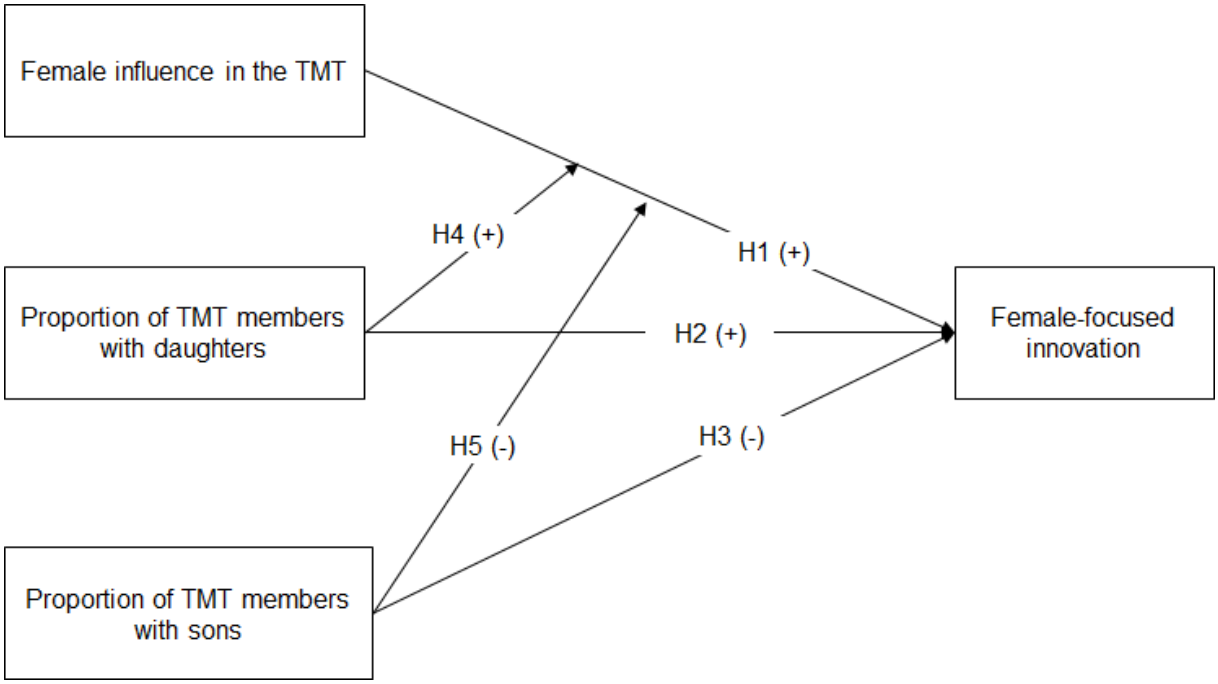

Figure 1.1: The Theoretical Framework

1.3.1 The relationship between female influence in the TMT and female-focused innovation

The attention-based view literature maintains that the issues to which members of the TMT allocate attention determine the type of innovation that firms introduce in the market (Li et al., 2013; Ocasio, 1997). Women leaders may bring distinct perspectives and insights related to the needs and preferences of female consumers (Dezsö and Ross, 2012), potentially suggesting a positive association between female influence in the TMT and female-focused innovation. However, whether these insights translate into innovation depends on how female executives engage with women-related issues—and whether they are willing to elevate these issues within the broader TMT agenda. Two contrasting perspectives offer different predictions regarding this relationship.

On one hand, the homophily perspective suggests that individuals are more likely to attend to, and advocate for, issues that reflect their own identity and lived experiences.

Within TMTs, this implies that as women's influence increases, so too does the salience of women-related concerns—heightening the likelihood that the firm will invest in innovations catering to female needs (Koning et al., 2021; Kozlowski et al., 2022). From an attention-based view, the more cognitive and strategic attention a topic receives, the more likely it is to be prioritized and resourced (Ocasio, 1997). Therefore, as female influence within the TMT increases, women can effectively channel attention toward gender-relevant domains, increasing female-focused innovation.

On the other hand, the queen bee hypothesis cautions that women in high-status roles may intentionally distance themselves from women-related issues to assimilate into male-dominated environments and gain legitimacy (Derks et al., 2016). In such contexts, advocating for female-focused innovation can be perceived as professionally risky, especially when gender is seen as a liability to success. To maintain their status, women in the TMT may avoid drawing attention to gender-specific needs—even when they are personally aware of them.

Thus, the queen bee perspective suggests that women in leadership positions may redirect the TMT's attention toward more traditionally "masculine" or neutral innovation domains.

Importantly, queen bee behavior is not an innate trait, but a situational response to organizational dynamics. It is more likely to emerge in environments where women are underrepresented, feel marginalized, or lack power to influence strategic outcomes (Corwin et al., 2022; Derks et al., 2016). When women hold limited influence, they may avoid championing gender-equity concerns in order to protect their position within the dominant power structure. In contrast, women with greater authority are more likely to challenge prevailing norms and legitimize women-related priorities. For instance, Corwin et al. (2022) find that women CEOs are less likely to appoint other women to their TMTs when facing strong external scrutiny or when their own power is constrained. However, as their power increases, this distancing effect diminishes.

These findings support the idea that queen bee behavior is context-dependent—shaped by the degree to which women feel empowered to deviate from prevailing gender norms. For this reason, we argue that as the female influence in the TMT increases, the homophily effect tends to outweigh the queen bee behavior, resulting in greater attention to women's needs and increased female-focused innovation. Thus, we hypothesize that:

Hypothesis 1. *Female influence in the TMT is positively associated with female-focused innovation*

1.3.2 The relationship between parenting and female-focused innovation

In addition to identity-driven mechanisms, situated attention theory suggests that individuals' attentional focus is shaped by their social environment (Ocasio, 1997). This perspective suggests that exposure to particular people, especially in emotionally significant settings, can influence what decision-makers perceive as strategically relevant or legitimate (Umashankar et al., 2022).

The daughter effect

The female socialization hypothesis builds on this logic by proposing that exposure to women in one's immediate environment, particularly in close and emotionally salient relationships, can shift parents' preferences and priorities in more gender-equitable directions (Glynn and Sen, 2015; Einiö et al., 2019; Washington, 2008; Warner and Steel, 1999). Past research shows that this shift is stronger for fathers than for mothers (Warner and Steel, 1999), partially because women were already in favor of gender-equality. In this regard, having daughters represents a powerful form of social exposure that may heighten male executives' attention to gender-related concerns, even if they do not share those concerns through demographic identity alone.

This effect can be explained through the mechanism of vicarious learning, whereby individuals acquire new understandings by observing and empathizing with others' experiences (Glynn and Sen, 2015). As male executives observe and engage with their daughters, they may come to appreciate the constraints and gender-based challenges women face—not through personal experience, but through relational proximity and emotional salience. This learning process reflects a key tenet of situated attention theory: that attention is not only a product of internal cognition, but is situationally cued and socially constructed. Daughters act as a repeated, salient stimulus that makes gender-relevant issues more cognitively accessible, emotionally resonant, and normatively legitimate. Applied to the innovation domain, these mechanisms suggest that executives with daughters may become more attuned to gender-specific market needs and more motivated to advocate for innovations that address them, hoping to improve the future lives of their daughters (Warner and Steel, 1999).

In this sense, daughters serve as a situational attentional trigger that reorients managerial cognition and decision-making toward gender-related concerns. Supporting this logic, empirical studies show that male political leaders with daughters are more likely to support women's rights (Glynn and Sen, 2015; Washington, 2008); CEOs with daughters invest more in gender-diverse hiring (Dasgupta et al., 2018); and male venture capitalists and founders with daughters are more inclined to gender-diverse hiring and promotion (Calder-Wang and Gompers, 2021; Wu et al., 2024).

Collectively, these findings suggest that having daughters acts as a situational, psychologically meaningful trigger that reorients managerial attention toward gender-related concerns, thereby increasing the firm's focus on female-focused innovation.

Thus, we hypothesize that:

Hypothesis 2. *There is a positive association between the proportion of TMT members with daughters and female-focused innovation.*

The son effect

While exposure to daughters may increase executives' awareness of gender equity concerns and reorient attention toward female-specific needs, the presence of sons may have the opposite effect, by failing to provide the contextual cues that make gender issues visible and salient. Thus, sons may reinforce traditional gender norms and sustain existing attentional schemas that deprioritize concerns linked to female consumers. Applied to the corporate domain, this means that executives with sons may become more attuned to traditionally male-dominated markets, innovation domains, or consumer segments, while placing less strategic weight on innovations that address women's needs. From a situated attention perspective, sons do not act as catalysts for attentional redirection, but as passive reinforcers of the status quo within the managerial cognitive environment, who leave intact the prevailing patterns of focus that often center on masculine-coded domains. In the absence of emotionally resonant, gender-relevant exposure—such as that provided by daughters—issues tied to women's experiences remain cognitively backgrounded, and may lack the perceived legitimacy or urgency needed to capture attention.

Empirical research supports this asymmetry. For example, past research reports that only fathers of daughters change their attitude and favor gender-equality policies, while men with sons have no incentive to challenge gender discrimination because their kids will not face this challenge (Warner and Steel, 1999). Similarly, male legislators (Washington, 2008), judges (Glynn and Sen, 2015), and venture capitalists (Cronqvist and Yu, 2017) consistently exhibit greater attention to gender-equity concerns when they have daughters, but not when they have sons. These findings suggest that sons maintain baseline attentional patterns, reinforcing the idea that, in the absence of female-relevant cues, gendered concerns remain outside of strategic focus.

This logic aligns with the broader tenets of selective attention theory: executives focus on domains they perceive as familiar, while overlooking those that lack immediacy or rel-

evance (Li et al., 2013; Ocasio, 2011; Umashankar et al., 2022). Sons, in this framework, fail to serve as disruptive or redirective cues, thereby allowing traditional attentional structures to persist. As a result, gender-equity concerns are less likely to be surfaced, championed, or translated into innovation priorities, even by women in the TMT, ultimately reducing the firm's attention to female-focused innovation. Thus, we hypothesize that:

Hypothesis 3. *There is a negative association between the proportion of TMT members with sons and female-focused innovation.*

1.3.3 The moderating effect of parenting

As outlined earlier, female executives in the TMT may bring distinct insights into gender-related needs and female consumers (Dezsö and Ross, 2012). However, consistent with the attention-based view, whether these insights translate into innovation depends not only on female influence, but also on the broader attentional context within the TMT. To capture this dynamic, we examine how the association between female influence and female-focused innovation varies based on the gender composition of TMT members' children.

The moderating effect of the proportion of TMT members with daughters

The presence of daughters among TMT members serves as a situated attentional cue, increasing the collective salience and perceived legitimacy of gender-related issues. When several executives have daughters, issues related to women's needs are more likely to be considered as strategically relevant, openly discussed, and integrated into decision-making. This attentional normalization reduces the perceived cost of advocating for female-focused innovation, particularly for female executives who might otherwise feel pressure to conform to dominant, male-centric norms. This mechanism is critical in light of the queen bee

effect, which suggests that women in male-dominated leadership settings may distance themselves from gender-equity issues to protect their status (Corwin et al., 2022; Derks et al., 2016). We argue that exposure to daughters among TMT members can buffer this effect. Indeed, when colleagues also exhibit attentional sensitivity to women's issues—as is the case when they have daughters—female executives may feel greater psychological safety and support to raise and act upon gendered insights. In such contexts, women leaders may be more empowered to direct innovation efforts in ways that reflect female consumers' needs.

In this way, daughters do more than influencing the attention of individual TMT members (H2); they help create a social environment where women leaders feel more legitimate in advocating for gender-equity concerns, amplifying the impact of female leadership on female-focused innovation. Thus, we propose that the positive association between female influence in the TMT and female-focused innovation is strengthened when a higher proportion of TMT members have daughters. We hypothesize that:

Hypothesis 4. *The proportion of TMT members with daughters positively moderates the relationship between female influence in the TMT and female-focused innovation.*

The moderating effect of the proportion of TMT with sons

As previously argued (H3), having sons does not introduce gender-relevant cues that reorient executive attention toward female needs. Instead, sons function as passive reinforcers of the status quo, sustaining existing attentional schemas that prioritize status-driven, male-oriented, or gender-neutral domains. From a situated attention perspective, sons fail to disrupt the traditional cognitive environment within TMTs (Ocasio, 1997). As a result, gender-equity concerns remain cognitively backgrounded—lacking the salience, urgency, or perceived legitimacy needed to drive strategic prioritization.

This attentional context matters for female executives within the TMT who seek to shape the firm's innovation agenda. When many members of the TMT have sons,

the social cues that heighten sensitivity to gender-equity issues are absent, making it more difficult for even influential women to redirect strategic focus toward female-oriented innovation. Moreover, this environment may activate the queen bee effect.

When women perceive that their gender-based insights are unsupported or illegitimate within the broader TMT, they may perceive a higher reputational cost in voicing such issues. To protect their status, female executives may engage in self-group distancing (Derks et al., 2015), weakening the link between female influence in the TMT and female-focused innovation.

In this way, sons do not simply fail to elevate attention to female-specific needs (H3); they may also dampen the effectiveness of female influence in the TMT by reinforcing an attentional context in which gender-equity concerns lack legitimacy. This increases the likelihood that women will adopt self-distancing behaviors and redirect innovation toward more traditional, male-oriented domains. Thus, we hypothesize that:

Hypothesis 5. *The proportion of TMT members with sons negatively moderates the relationship between female influence in the TMT and female-focused innovation.*

1.4 Method

1.4.1 Sample

We construct a panel dataset spanning 77 quarters from 2000 to 2019 to examine the relationship between executives' personal backgrounds and firms' innovation focus. The dataset integrates multiple sources to capture firm characteristics, top management composition, and innovation outcomes.

Patent data are obtained from the United States Patent and Trademark Office and linked to corporate entities using the KPSS Technological Innovation, Resource Allocation, and Growth dataset, which matches patent assignees to firm identifiers. Firm finan-

cial and governance data come from Compustat, BoardEx, and ExecuComp, providing information on company fundamentals, executive demographics, and team composition. We further incorporate consumer spending and behavioral data from L2 data, which are linked to executives using a crosswalk developed by Raffiee et al. (2022). This linkage ensures consistent identification of top managers across data sources. Appendix A1 provides a detailed walkthrough of the data integration process.

We focus on firms operating in the pharmaceutical industry (SIC 283). This setting is appropriate for two reasons. First, pharmaceuticals represent a knowledge-intensive sector in which innovation is central to firm success (Chandy et al., 2006; Sorescu et al., 2003), paralleling other industries frequently examined in studies of technological innovation (Chandy and Tellis, 2000). Second, prior research documents systematic gender disparities in pharmaceutical innovation, with firms underinvesting in innovations targeting female health needs (Koning et al., 2021). Given the high rate of patenting in this industry to secure intellectual property (Prabhu et al., 2005), we employ patents as our measure of innovation. To best capture the timing of firms' innovation decisions and align them with executive characteristics, we rely on the patent *application* date rather than the grant date.

The final dataset includes 5,545 firm–quarter observations across 148 unique firms. We restrict the sample to the pre-pandemic period to avoid potential distortions introduced by COVID-19–related disruptions in innovation and organizational behavior.

1.4.2 Measures

Dependent variable: Female-focused innovation

To quantify the extent to which firms focus their innovation efforts on female-specific needs, we construct a measure based on patent text analysis, following Koning et al. (2020). Using the KPSS match between patents and firms, we analyze the titles and

abstracts of each patent.

We employ a classification model based on a Generative Pre-trained Transformer (GPT), developed by OpenAI (Devlin et al., 2018; Radford et al., 2018), to infer the intended demographic orientation of each patent. Each patent is categorized as addressing needs specific to females, males, or neither. We assign numeric scores to these categories, enabling us to capture not only whether a firm targets gender-specific needs, but also in which direction.

Firm-level female focus is constructed as the average orientation score across all patents filed in a given quarter:

$$\text{Female-Focused Innovation}_{it} = \frac{1}{P_{it}} \sum_{i \in P_{it}} \text{Score}_j, \tag{1.1}$$

where P_{it} is the number of patents filed by firm i in quarter t, and $\text{Score}_j \in \{-1, 0, +1\}$ denotes the demographic orientation assigned to patent j. The resulting measure ranges from -1 (entirely male-focused) to $+1$ (entirely female-focused), with values near zero indicating gender neutrality.

Appendix A2 provides representative examples of patents classified as female- and male-focused, together with Medical Subject Headings (MeSH) terms extracted from their abstracts, offering face validity for the classification.

Independent variables

We define the top management team (TMT) as all C-level executives, including the CEO, as well as executives at or above the rank of vice president or senior vice president, following Varma et al. (2023). BoardEx provides detailed start and end dates for each executive's tenure, enabling us to construct a dynamic quarterly panel of TMT composition at the firm level.

Female influence in the TMT. We adopt the composite index developed by Srivastava

et al. (2023), which incorporates four indicators: (1) the proportion of female executives in the TMT, (2) the hierarchical rank of the highest-ranking female executive, (3) the cumulative ranking scores of all female executives, and (4) the number of roles held by female executives. These indicators jointly capture both representation and authority. Following Srivastava et al. (2023), we combine them using factor analysis and rescale the resulting index to range between 0 and 1 for comparability over time and across firms.

Proportion of TMT members with daughters/sons. Data on executives' children come from L2 data, which provides cross-sectional information on child gender and age bins. Although L2 data is not longitudinal, we construct quarterly estimates by using the midpoint of each age bin to approximate each child's birth year. For each TMT member, we infer the periods during which they likely had at least one daughter or son. Aggregating across executives yields firm-quarter level measures of the proportion of TMT members with at least one daughter and the proportion with at least one son. This approach introduces plausible time variation into family-structure variables and allows us to examine how TMT exposure to daughters or sons relates to firms' female-focused innovation.

Control variables

At the firm level, we control for the total number of patents filed in a given quarter, which captures baseline levels of innovative activity and prevents our female-focused innovation measure from conflating shifts in gender orientation with changes in overall innovation volume. We also include a measure of *Marketing Power in the TMT* similar to Feng et al. (2015), which reflects the influence of marketing executives within the leadership team. Marketing-led teams may be more attuned to consumer needs, potentially increasing support for gender-inclusive innovation; however, they may also favor broad-market strategies that deprioritize gender-specific product development. Given these competing theoretical expectations, we remain agnostic regarding the direction of this effect.

At the individual level, we incorporate several aggregated characteristics of the TMT to account for executives' personal values and demographic backgrounds. Political ideology is captured by the proportions of executives registered as *Democrats* and *Conservatives*, reflecting well-established links between political orientation and attitudes toward gender equity (Warner and Steel, 1999). We further control for the average number of children in executives' households and the proportion of executives belonging to *Christian families*, as family structure and religious affiliation are known correlates of gender-related beliefs and preferences. Finally, we include the proportion of executives who engage in *women-related spending*, which may signal personal familiarity with or attentiveness to female-oriented product categories. Period fixed effects are included in all specifications to account for quarter-specific shocks affecting the entire industry. Table 1.1 reports an overview of all variables used in the analysis, including their data sources and construction procedures.

1.4.3 Estimation strategy

We test our theoretical framework with the following random effects panel regression model:

$$\begin{aligned}\text{Female-Focused Innovation}_{it+1} = \beta_0 &+ \beta_1 \text{FITMT}_{it} + \beta_2 \text{Prop. Daughters}_{it} + \beta_3 \text{Prop. Sons}_{it} \\ &+ \beta_4(\text{FITMT}_{it} \times \text{Prop. Daughters}_{it}) + \beta_5(\text{FITMT}_{it} \times \text{Prop. Sons}_{it}) \\ &+ \sum_{p=1}^{P} \gamma_p \text{Period}_p + \sum_{k=1}^{6} \delta_k \text{Control}_k + u_i + \varepsilon_{it} \qquad (1.2)\end{aligned}$$

where the subscripts i and t refer to firm and quarter, respectively; *FITMT* denotes female influence in the TMT, and Prop. Daughters and Prop. Sons represent the proportion of TMT members with daughters and sons, respectively. The interaction terms allow us to test hypotheses H4 and H5. The vector *Control* includes the eight firm- and team-level controls described above. Period fixed effects account for time-specific shocks

Variable	Source	Construction / Operationalization
Female-Focused Innovation (DV)	USPTO	GPT-based scoring of patent titles and abstracts (1 = male-focused; +1 = female-focused). Scores averaged across all patents applied for by the firm in each quarter.
Female Influence in the TMT (FITMT)	BoardEx	Composite index following Srivastava et al. (2023), combining: (1) proportion of female executives, (2) highest female hierarchical rank, (3) cumulative female rank score, (4) number of roles held by female executives. Indicators combined via factor analysis and normalized to range [0,1].
Proportion of TMT Members with Daughters	L2 data	Child gender and age-bin information used to infer approximate birth years. Quarterly indicator for executives with $\geq$1 daughter. Aggregated to firm-quarter proportion.
Proportion of TMT Members with Sons	L2 data	Constructed analogously to the daughters measure, indicating executives with $\geq$1 son. Aggregated to firm-quarter proportion.
Marketing Power in the TMT	BoardEx	Weighted presence of marketing-related titles in the TMT, inspired by Feng et al. (2015). Centered and scaled.
Democrats	L2 data	Proportion of TMT members registered as Democrats. Baseline group in estimation = politically neutral members.
Conservatives	L2 data	Proportion of TMT members registered as Conservatives. Baseline group = politically neutral members.
Number of Children	L2 data	Number of children per TMT member, average aggregated to the firm-quarter level.
Christian Families	L2 data	Proportion of TMT members belonging to Christian religious households.
Women-Related Spending	L2 data	Proportion of TMT members with spending patterns in female-oriented product and service categories.

Table 1.1: Variables Sources and Construction

affecting all firms.

We forward-lag the dependent variable by one quarter to ensure temporal ordering and to reduce concerns about reverse causality, allowing us to examine whether current TMT characteristics predict future female-focused innovation.

Addressing endogeneity concern of female influence in the TMT

Female influence in the TMT may be endogenous (Srivastava et al., 2023). In our empirical model, we include period fixed effects to account for time-specific shocks and control for a range of observable firm- and TMT-level characteristics that vary over time. However, unobserved time-varying factors, such as internal cultural or strategic priority shifts not captured in our data, may still influence a firm's decision to appoint or promote women to top management roles. These omitted variables could bias our estimates if they are also correlated with innovation outcomes.

To address this potential endogeneity, we implement a control function approach, which has been widely used to correct for endogeneity in both linear and nonlinear settings (Petrin and Train, 2010). This method involves estimating an auxiliary regression in which female influence in the TMT is the dependent variable, using the following specification:

$$\begin{aligned} \text{FITMT}_{it} = \alpha_0 + \alpha_1 \text{WPFITMT}_{it} + \alpha_2 \text{Prop. Daughters}_{it} + \alpha_3 \text{Prop. Sons}_{it} \\ + \sum_{p=1}^{P} \gamma_p \text{Period}_p + \sum_{k=1}^{8} \delta_k \text{Control}_k + \mu_i + \varepsilon_{it} \end{aligned} \tag{1.3}$$

Specifically, we regress our endogenous variable—female influence in the TMT (*FITMT*)—on an instrumental variable (*WPFITMT*) along with the two other independent variables (the proportions of TMT members with daughters and with sons) and the same control variables included in Equation 1.2. We include firm fixed effects (μ_i) and correct standard errors for autocorrelation using the `xtregar` procedure. We then plug the estimated

control function residuals from Equation 1.3 into the main models in Equation 1.2. We use `xtregar` rather than `xtreg` because our panel exhibits first-order serial correlation in the idiosyncratic error term ε_{it}; failing to account for autocorrelation can lead to biased standard errors and unreliable inference. A Wooldridge test confirms the presence of AR(1) serial correlation ($F(1, 145) = 65554.57$, $p = 0.000$), justifying the use of xtregar for our fixed-effects specification. A Hausman specification test comparing fixed-effects and random-effects models does not reject the null of no systematic coefficient differences ($\chi^2(10) = 9.07$, $p = 0.526$) supporting the use of a random-effects specification.

Following Srivastava et al. (2023), we use as our instrumental variable the weighted average of female influence in the TMT among geographically proximate, industry-excluded firms ($WPFITMT$). For each firm-quarter, peers are defined as firms headquartered in the same U.S. State but operating in different industries, allowing us to capture localized institutional or normative pressures for gender inclusion in leadership without conflating industry-specific strategic dynamics. To increase comparability, peer firms are weighted based on financial similarity using multidimensional scaling across key variables (assets, liabilities, sales, R&D, and market value). This approach ensures that influence is drawn more heavily from financially similar firms. We argue that while gender inclusion norms may cluster geographically, these local pressures are unlikely to directly influence how a pharmaceutical firm allocates innovation resources, given the global nature of drug development and sales. As shown in Table 1.3, Model 0, the instrument is strongly and significantly associated with female influence in the TMT ($\alpha_1 = 1.262$, $p = 0.02$), satisfying the relevance criterion. By excluding same-industry firms, we also mitigate concerns about the exclusion restriction, which are particularly salient in innovation contexts where industry peers can affect strategic decisions through benchmarking or imitation (Varma et al., 2023). In contrast, geographically proximate firms in unrelated industries are unlikely to shape a focal firm's innovation agenda except through shared social or institutional norms that influence TMT composition.

Accounting for sample selection: Heckman correction

Because our dependent variable, female-focused innovation, is only observed for firms that file at least one patent in a given quarter, our estimation sample may be subject to non-random selection. Firms that innovate may systematically differ from those that do not, potentially biasing our regression estimates. To address this issue, we implement a Heckman (Heckman, 1976, 1979) selection correction.

We estimate a Probit model predicting whether a firm engages in innovation activity:

$$\text{Innovator}_{it} = \begin{cases} 1 & \text{if Total Patents}_{it} > 0, \\ 0 & \text{otherwise,} \end{cases} \tag{1.4}$$

$$\Pr(\text{Innovator}_{it} = 1) = \Phi\left(\begin{array}{l} \theta_0 + \theta_1 \text{Innovator}_{i,t-1} + \theta_2 \ln(\text{Sales}_{it}) \\ + \theta_3 \text{Marketing Power}_{it} + \theta_4 \text{Team Size}_{it} \\ + \theta_5 \text{State}_i + \sum_{p=1}^{P} \eta_p \text{Period}_p \end{array}\right). \tag{1.5}$$

where $\Phi(\cdot)$ is the cumulative normal distribution function.

Using the fitted values $\hat{x}\theta$, we compute the inverse Mills ratio,

$$\lambda_{it} = \frac{\phi(\hat{x}\theta)}{\Phi(\hat{x}\theta)}, \tag{1.6}$$

where $\phi(\cdot)$ denotes the standard normal density function.

In the second step, we augment the main outcome equation (Equation 1.2) by including the inverse Mills ratio λ_{it} as an additional regressor. This procedure corrects for potential selection bias arising from the fact that female-focused innovation is only observed for the subset of firm-quarters in which patenting occurs.

Addressing endogeneity concerns for gender composition of TMT members' children

Our identification strategy rests on the assumption that the gender of TMT members' children is effectively exogenous to the prior strategic preferences and characteristics of TMT members. This follows the logic of Washington (2008) and Glynn and Sen (2015), who treat child gender as a quasi-random assignment that enables causal inference. Once an individual becomes a parent, the likelihood of having a daughter versus a son is assumed to be beyond their control and uncorrelated with unobserved traits such as innovation preferences or gender-related attitudes. This exogeneity allows us to interpret the observed associations between proportion of TMT members with daughters/sons and female-focused innovation as plausibly causal.

A potential concern is that child gender may not be entirely random within our executive sample. For instance, TMT members may have engaged in sex selection or other behaviors correlated with unobserved attitudes or preferences. However, the average age of TMT members in our dataset is approximately 50, suggesting that most children were likely born before fetal sex-selection technologies became widely available or accessible. Additionally, so-called "natural" methods of influencing child gender, such as timing of intercourse, have not been shown to produce reliable effects (Wilcox et al., 1995). While we cannot fully rule out cases involving adoption or step-parenting, such arrangements are expected to be uncommon in this population and unlikely to generate systematic bias in our estimates.

Another potential concern is selection into top management roles based on personal characteristics correlated with child gender, such as the possibility that boards may implicitly favor or disfavor candidates based on family composition. However, there is no theoretical or empirical evidence to suggest that such selection mechanisms operate at the executive level. Moreover, our analysis includes a comprehensive set of controls, such as

the average number of children and key ideological and behavioral traits of TMT members, that account for variation in both family structure and personal values, further mitigating concerns about omitted variable bias.

1.5 Results

1.5.1 Main analyses

We mean-center the continuous variables before creating the interaction terms to reduce multicollinearity (Aiken et al., 1991). We report our results in Table 1.3.

Model 1 examines the direct associations between female influence in the TMT and female-focused innovation. The coefficient for female influence in the TMT is negative and statistically significant ($\beta = -0.0156$, $p = 0.000$), indicating that, contrary to our expectations in H1, greater female influence within the TMT is associated with lower levels of female-focused innovation. Economically, a one standard deviation increase in female influence in the TMT corresponds to a 0.41 percentage point decrease in female-focused innovation. This unexpected finding may reflect the presence of queen bee dynamics, particularly when female leaders are embedded in male-dominated environments with limited legitimacy to advocate for gender-related issues, as in the case of firms in our sample.

Turning to H2 and H3, we find that the proportion of TMT members with daughters is positively associated with female-focused innovation, although not statistically significant in Model 1 ($\beta = 0.00389$, $p = 0.113$). In contrast, the proportion of TMT members with sons is negatively associated with female-focused innovation and statistically significant ($\beta = -0.0079$, $p = 0.002$), supporting H3. Because both parental-gender variables range from 0 to 1, these coefficients represent meaningful shifts: moving from a TMT in which none of the executives have sons to one in which all executives have at least one son corresponds to a 0.79 percentage point decrease in the female-focus score.

Variables	Selection
Lagged innovator	1.1673***
	(0.0420)
Log sales	0.1198***
	(0.0091)
Marketing power in TMT	-0.3189
	(0.2723)
Team size	0.0360***
	(0.0116)
State fixed effects	Yes
Period fixed effects	Yes
Constant	0.2459
	(0.2282)
Observations	5,213
Log-likelihood	-2636.2705
LR chi^2	1947.79
Pseudo R^2	0.2698

Note: *** p<0.01, ** p<0.05, * p<0.1. Standard errors in parentheses.

Table 1.2: Heckman Selection Equation

Model 2 introduces interaction terms to test whether the relationship between female influence in the TMT and female-focused innovation is moderated by the parental gender composition of the TMT. Specifically, we examine whether the effect of female influence in the TMT is conditioned by the proportion of executives with daughters (H4) or sons (H5). The interaction between female influence in the TMT and the proportion of TMT members with daughters is positive ($\beta = 0.048$, $p = 0.00$), while the interaction with the proportion of TMT members with sons is negative ($\beta = -0.055$, $p = 0.00$). Thus, we find support for H4 and H5, respectively. These results indicate that the association between female influence in the TMT and female-focused innovation is not uniform, but it depends critically on the gender of the TMT members' children.

To better interpret these interactions, we apply the Johnson–Neyman technique (Johnson and Neyman, 1936) and plot the conditional marginal effect of female influence across the empirical distribution of each moderator (Figure 1.2). Panel A shows that the relationship between female influence in the TMT and female-focused innovation is negative

Variables	First Stage Model 0	Second Stage Model 1	Second Stage Model 2
Weighted average of peer firms' FITMT (WPFITMT)	1.2623**		
	(0.5408)		
Female Influence in the TMT (FITMT)		-0.0156***	-0.0134***
		(0.0034)	(0.0034)
Proportion TMT members with daughters	-0.0631***	0.0039	0.0055**
	(0.0140)	(0.0026)	(0.0026)
Proportion TMT members with sons	0.0490***	-0.0079***	-0.0105***
	(0.0127)	(0.0025)	(0.0026)
FITMT*Proportion TMT members with daughters			0.0482***
			(0.0098)
FITMT*Proportion TMT members with sons			-0.0558***
			(0.0114)
Marketing Power in the TMT	0.1357***	0.0025	0.0035
	(0.0236)	(0.0088)	(0.0088)
Total patents	0.0002	0.0001	0.0001
	(0.0002)	(0.0001)	(0.0001)
Democrats	-0.0441***	-0.0002	-0.0006
	(0.0104)	(0.0022)	(0.0022)
Conservative	-0.0437***	-0.0005	-0.0014
	(0.0084)	(0.0019)	(0.0019)
Number of children households	0.0116*	0.0013	0.0010
	(0.0068)	(0.0013)	(0.0013)
Christian Families	-0.0116	-0.0002	-0.0008
	(0.0158)	(0.0029)	(0.0029)
Women Product Expenses	0.0263**	0.0005	0.0007
	(0.0113)	(0.0020)	(0.0020)
Control function residual		0.0183***	0.0187***
		(0.0051)	(0.0051)
Mills ratio		0.0043***	0.0039**
		(0.0016)	(0.0016)
Constant	0.0713***	0.0030	0.0038
	(0.0229)	(0.0087)	(0.0087)
Observations	5,545	4,917	4,917
Number of firms	148	137	137
Period Fixed Effects	Yes	Yes	Yes

Note: *** p<0.01, ** p<0.05, * p<0.1. Standard errors in parentheses.

Table 1.3: Panel Analysis

when the proportion of TMT members with daughters is below 0.33, a condition that holds in 53.7% of the sample. Conversely, when the proportion exceeds 0.86 the relationship becomes significantly positive, covering 19% of the sample. This suggests that a critical mass of TMT members with daughters is required to reverse the negative effect of female influence. Panel B illustrates the moderating effect of sons. The relationship between female influence in the TMT and female-focused innovation becomes significantly negative when the proportion of TMT members with sons exceeds 0.05, a threshold met in 44.35% of the sample. This suggests that only a small degree of TMT members with sons is sufficient to suppress any positive effect of female influence, highlighting an asymmetry in how child gender shapes attentional dynamics. Together, these results underscore a key insight: female influence in the TMT alone is not enough to increase attention to female-focused innovation. Instead, the social context of the TMT—particularly the gendered parenting experiences of its members—can either amplify or dampen the strategic salience of women's needs. Daughters increase sensitivity and can unlock the positive potential of female influence, while sons appear to reinforce traditional attentional patterns, limiting the expression of that innovation in gendered innovation choices.

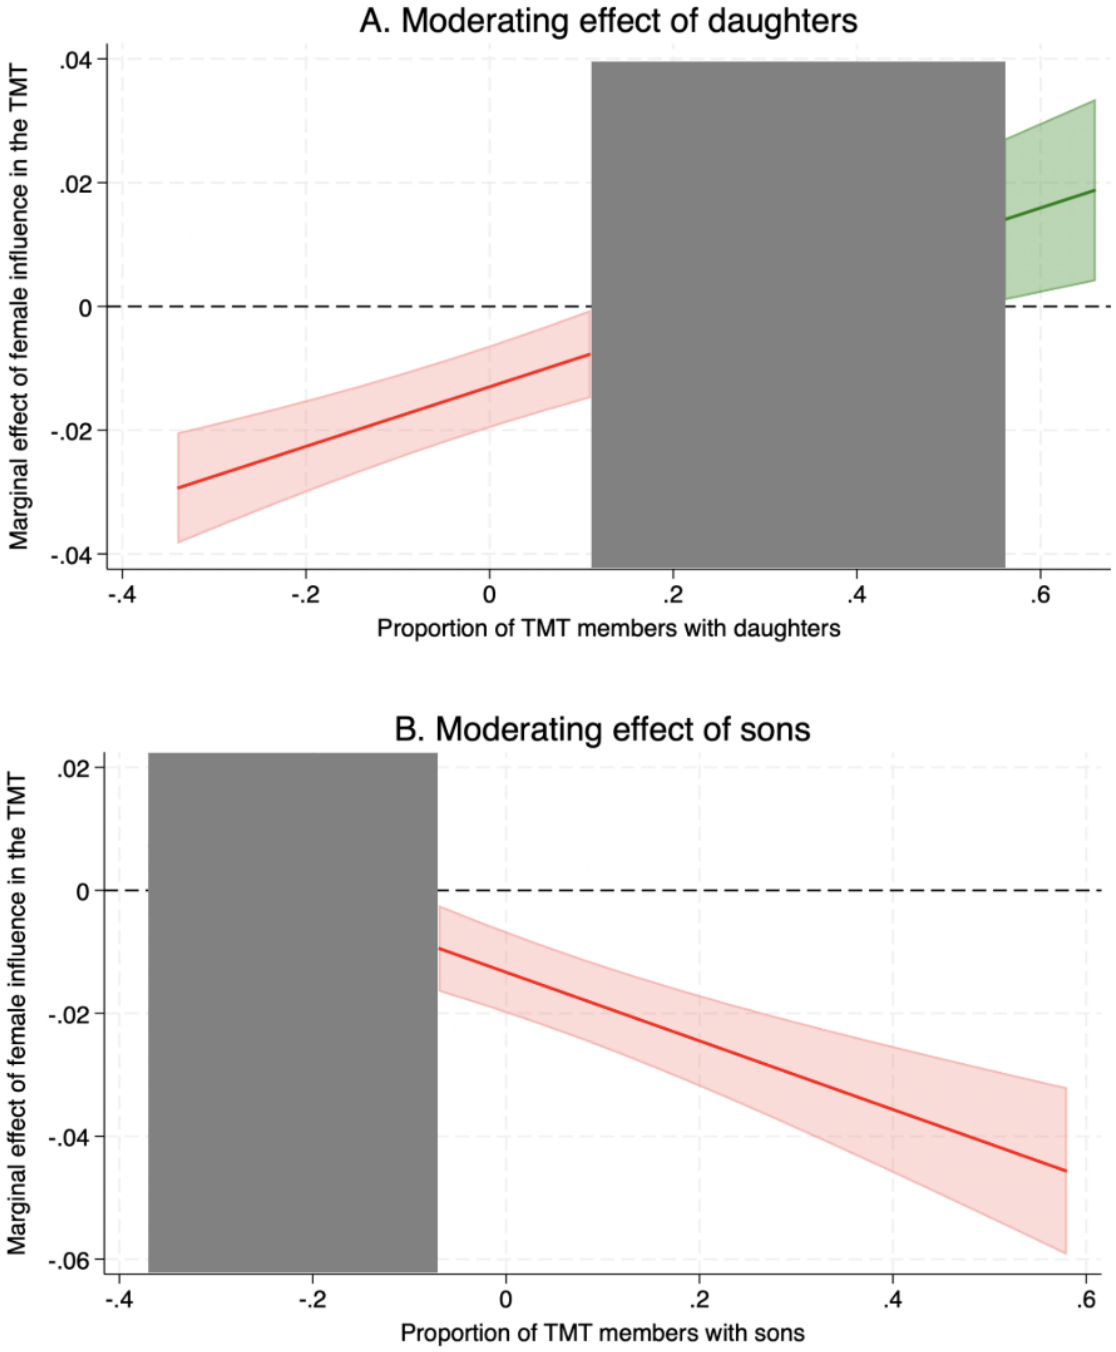

Figure 1.2: Interaction effects of TMT parental gender composition and FITMT on female-focused innovation

1.5.2 Robustness analyses

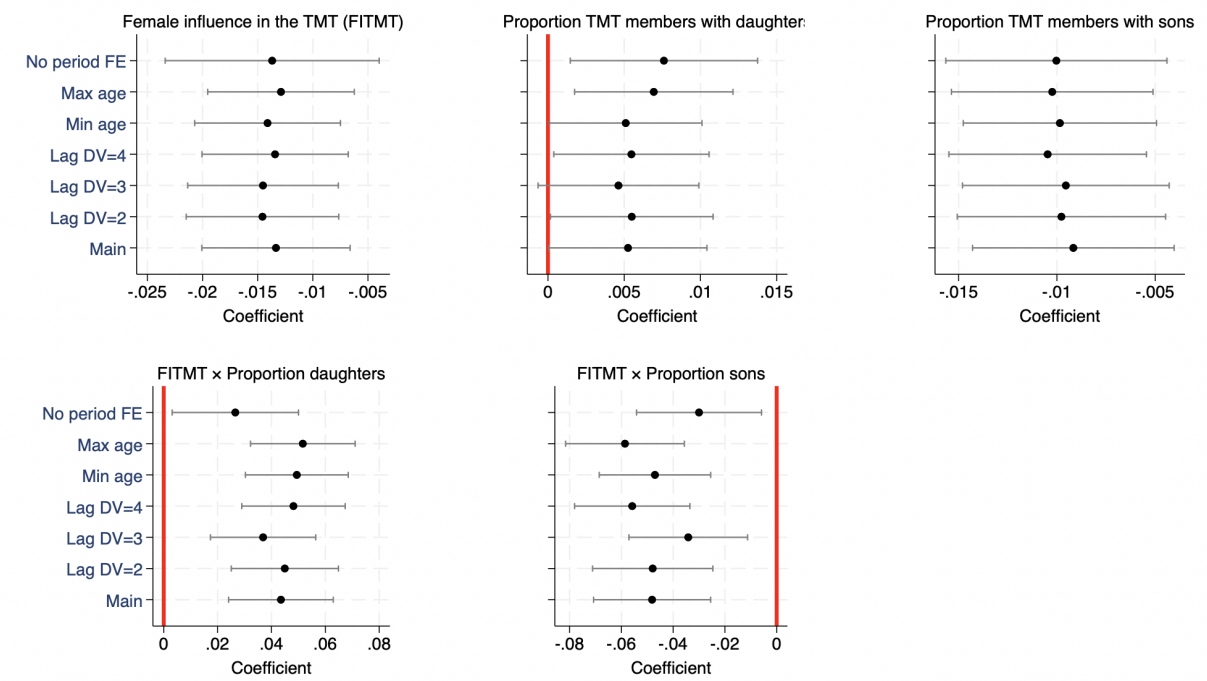

Figure 1.3: Specification Maps

Figure 1.3 presents specification maps that summarize the results of a series of robustness checks. First, we re-estimate our models using alternative forward-lags of the dependent variable, ranging from two to four quarters. Across all specifications, the sign and statistical significance of the key coefficients remain unchanged. Second, we employ alternative measures of the independent variables, capturing the proportion of TMT members with daughters or sons. We test alternative constructions by replacing the average estimated date of birth with the minimum and maximum plausible dates, derived from the lower and upper bounds of each child's reported age bin. These variations reflect more conservative or generous assumptions about when the eldest child was present in the household. Again, the signs and the significance of the coefficients remain consistent. Third, we estimate models excluding period fixed effects, and the coefficients do not change. Collectively, these analyses suggest that our results are robust to lag structure,

measurement of key predictors, or model specification.

1.6 Discussion

This study set out to examine how gender dynamics within TMT shape the firm's innovation focus, particularly when it comes to addressing the needs of female consumers. While prior work has emphasized the benefits of gender diversity in leadership (Srivastava et al., 2023), it remained unclear whether female presence in the TMT reliably translates into strategic prioritization of women-related concerns. We addressed this gap by integrating identity-based and context-based mechanisms through the lens of the attention-based view of the firm.

1.6.1 Theoretical contributions

Our findings contribute to three streams of research: the role of women in leadership positions (Dezsö and Ross, 2012; Post and Byron, 2015; Srivastava et al., 2023), the attention-based view (Ocasio, 1997), and the gendered innovation literature (Koning et al., 2021; Kozlowski et al., 2022; Lyngsie and Foss, 2017). We now describe these contributions in more detail.

Female leadership and the gender-innovation paradox. Our study contributes to the growing literature on the role of women in the TMT, and more in general in firm leadership positions, by advancing the conversation from questions of representation to questions of attentional focus. While prior research has established that gender diversity can positively affect firm outcomes—such as innovation (Torchia et al., 2011), customer orientation (Srivastava et al., 2023), and financial performance (Dezsö and Ross, 2012; Post and Byron, 2015)—we show that when it comes to female-focused innovation, these benefits hinge not just on women's presence or influence, but on the contextual support to focus a firm's innovation on female needs.

Our findings challenge a widely held assumption in this literature: that increasing female presence or influence will inherently lead to more inclusive, gender-equitable firm outcomes (Graves Jr et al., 2022; Iyer, 2020; Mickey and Smith-Doerr, 2022). These studies suggest that women bring distinct perspectives and stakeholder sensitivities that, when present in sufficient numbers or at high ranks, enhance firm responsiveness to underserved customers, such as women. However, our findings complicate this narrative. We show that, contrary to expectations, greater female influence in the TMT may actually reduce female-focused innovation. This counterintuitive result suggests that under certain conditions, particularly in male-dominated environments such as current TMTs, female leaders may distance themselves from gender-related issues rather than champion them. In doing so, our study empirically supports concerns raised in the literature on queen bee behavior (Derks et al., 2015), showing how identity threat and legitimacy concerns can suppress the very advocacy that increased representation is expected to promote. This insight pushes the literature forward by emphasizing that the impact of gender diversity is conditional and context-dependent, not automatic or uniformly positive.

Extending situated attention through parenting. Ocasio (1997, p. 204) considers situational attention the "central contribution" of his theory. However, it received scant attention in the literature. Our study contributes to the attention-based view by empirically unpacking one of its core, yet underexplored, tenets: situated attention. While foundational work by Ocasio (1997, 2011) theorizes that attention is shaped not only by individual traits but also by social and organizational context, much of the empirical literature has focused on internal attention structures such as organizational architecture or structural positions (see Brielmaier et al., 2023 for a review). In contrast, our work extends the attention-based toward a deeper integration of external, identity-relevant contexts that shape how decision-makers attend to and act on strategic issues.

Specifically, we investigate parenting experience. We show that parenting daughters serves as an emotionally salient, identity-relevant cue that increases executives' attention

to gender-related concerns, thereby fostering female-focused innovation. Conversely, parenting sons has the opposite effect: it reduces attentional sensitivity to these concerns and steers the firm away from female-focused innovation. In this way, we empirically demonstrate that personal family structures can produce differential and even opposing shifts in strategic attention, depending on the gender of the child. These findings offer a novel operationalization of the attention-based view's concept of "environmental embeddedness": not as broad industry or competitive conditions (Brielmaier and Friesl, 2023), but as intimate, social environments that feed back into the strategic focus of the firm via the attentional schemas of its leaders. These insights allow us to move beyond the traditional organizational boundaries of the attention-based view research and demonstrate how relationships in the private sphere influence collective resource allocation in the corporate domain.

By incorporating both demographic identity (i.e., female executives in the TMT) and situational reinforcement (parenting daughters or sons), we bridge the attention-based view with insights from upper-echelons theory and social psychology, offering a more complete account of how gendered attention emerges in firm decision-making. While upper echelons theory has long emphasized the role of identity in shaping managerial cognition (Hambrick and Mason, 1984), we draw on social psychological theories of vicarious learning to explain how emotionally meaningful, identity-relevant relationships can reorient attentional focus, even among those who do not share the identity in question. Crucially, our findings show that the influence of women in the TMT does not guarantee increased attention to female-focused innovation unless it occurs within a supportive attentional environment that makes gender-relevant issues salient and legitimate. This finding represents a key extension of the attention-based view: attention is not just situated, but it is *socially* situated, and often *emotionally* cued. Our study thus advances a more behaviorally realistic and sociologically grounded account of how managerial attention is formed and focused.

Extending the gender-innovation research upstream. Recent work has shown that who invents shapes what gets invented. Studies like Koning et al. (2020), Kozlowski et al. (2022), and Gagliardi and Novelli (2024) show that women inventors and entrepreneurs are significantly more likely to focus on women's health and needs, resulting in a "gender gap" in innovation content driven by the underrepresentation of women among inventors. This phenomenon is often explained by the link between an innovator's identity and their awareness of underserved needs in the marketplace. Yet, they predominantly focus on inventors, scientists, or user-entrepreneurs, offering limited insight into the upstream organizational processes that determine which ideas receive resources and strategic support. Our study moves one step up the organizational hierarchy by examining TMTs —the decision-makers who guide firm-level innovation strategy. As Koning et al. (2020) note, the gender gap in commercialization contributes to the gender gap in invention outcomes. We extend this logic by asking: who decides what innovation gets developed in the first place? By focusing on female influence and parenting experiences of TMT members, we explore how leadership dynamics—not just inventor identity—shape the attentional focus and innovation at the firm level. In doing so, our study adds an important leadership lens to the growing literature on gender and innovation, highlighting the pivotal role of upstream decision-making in the production of inclusive innovation outcomes.

1.6.2 Managerial implications

Our findings carry an important message for marketing leaders and corporate decision-makers seeking to better serve an economically vital, yet underserved market segment, such as female consumers. While many firms have embraced gender diversity in leadership as a pathway to more female-focused innovation, our results suggest that representation alone is not enough. In organizations where gender-related concerns are perceived as illegitimate or risky to champion, even highly placed women may feel pressure to conform to dominant norms, thereby muting their potential strategic influence. This insight calls

for a shift from focusing solely on "who is at the table" to how supported and empowered those individuals are once they are there. Firms must cultivate cultures where gender-specific needs are not only visible but also strategically prioritized and rewarded.

Second, while our findings reveal that female influence in the TMT is not consistently associated with more female-focused innovation, they also point to a promising lever: executives' parenting experiences. Specifically, TMT members with daughters exhibit greater attentional sensitivity to gender-related concerns, which correlates with increased female-focused innovation. This finding suggests a practical pathway for firms seeking to better serve female consumers, even when structural constraints limit the immediate appointment of more women to leadership roles. By recognizing and leveraging the experiential diversity of male leaders—particularly those with daughters—firms can tap into authentic, emotionally grounded insights that enhance strategic focus on underserved markets. Appointment and promotion decisions can take such lived experiences into account. In doing so, firms can expand their tools for fostering female-focused innovation beyond traditional gender metrics.

Third, and more broadly, our findings speak to ongoing efforts to increase diversity in corporate leadership, not only in terms of gender but also race, ethnicity, and other historically marginalized identities, with the goal of improving equity and responsiveness to underserved consumer segments. These initiatives rest on the assumption that increasing representation at the top will naturally lead to more inclusive organizational priorities and outcomes. However, our results suggest that this link is more conditional than commonly assumed. When marginalized leaders operate in environments where their identities are devalued or unsupported, their ability—or willingness—to advocate for their communities' needs may be constrained. While our study centers on gender, similar dynamics likely affect racial or ethnic minority executives, particularly in predominantly white or non-inclusive leadership settings. These insights offer a valuable lens for firms and policy makers seeking to move beyond surface-level diversity and toward more structurally

inclusive outcomes.

1.6.3 Limitations and Directions for Future Research

Our study has several limitations that open avenues for future research. First, we approximate executives' parenting experiences using cross-sectional data on child gender and age bins. While this allows us to introduce plausible time variation, our measure remains an estimate and does not capture the full complexity or depth of parent–child relationships. Future research could benefit from more granular or self-reported data on family dynamics and values.

Second, our context is limited to the pharmaceutical industry, which, while well-suited to studying female-focused innovation, may not generalize to other sectors with different innovation cycles or gender dynamics. Studies in other industries—such as technology, finance, or consumer goods—could help validate and extend our findings.

Third, our measure of female-focused innovation, though novel and rigorous, focuses on patent data and may not capture other important forms of inclusion, such as product design, marketing strategies, or service delivery. Broader measures of innovation equity warrant further exploration.

Finally, we focus on one specific underserved segment: female consumers. Future research should explore whether and how leadership influence, emotional connection, and contextual support shape the ability of other underrepresented groups—such as Black, LGBTQ+, or disabled individuals—to increase innovation addressing the needs of the communities they are presumed to represent.

A1 Appendix: Construction of the Final Dataset

To construct the dataset used in this study, we integrate multiple sources covering firm characteristics, top management composition, and innovation outcomes. The dataset is built through the following steps.

First, we extract all firms operating in the pharmaceutical industry (SIC 283) from Compustat. Next, we identify top management team (TMT) members using BoardEx, selecting all executives at the level of vice president and above while excluding board members. Since BoardEx does not provide gender information for all executives, one of the authors manually verified and supplemented missing data using external sources such as LinkedIn and Bloomberg. When gender was not explicitly stated, it was inferred based on photographs or gendered language in professional biographies.

Innovation data is obtained from the United States Patent and Trademark Office (USPTO) and PatentsView.com, covering all patents filed between 2000 Q1 and 2019 Q3. These patents are matched to firm-level data using the KPSS Technological Innovation, Resource Allocation, Growth dataset, which links patent assignees to their respective firm GVKEYs. The patent data is then aggregated at the firm-quarter level.

To incorporate executive characteristics, we link the firm-level dataset to ExecuComp using a GVKEY-company ID crosswalk. Next, we match executives from ExecuComp to the L2 Data dataset using the DirectorID from ExecuComp and the VoterID in L2 Data, exploiting the crosswalk developed by Raffiee et al. (2022). This allows us to integrate individual-level demographic and behavioral characteristics of executives. All executive-level information is aggregated at the firm-quarter level to maintain consistency across datasets.

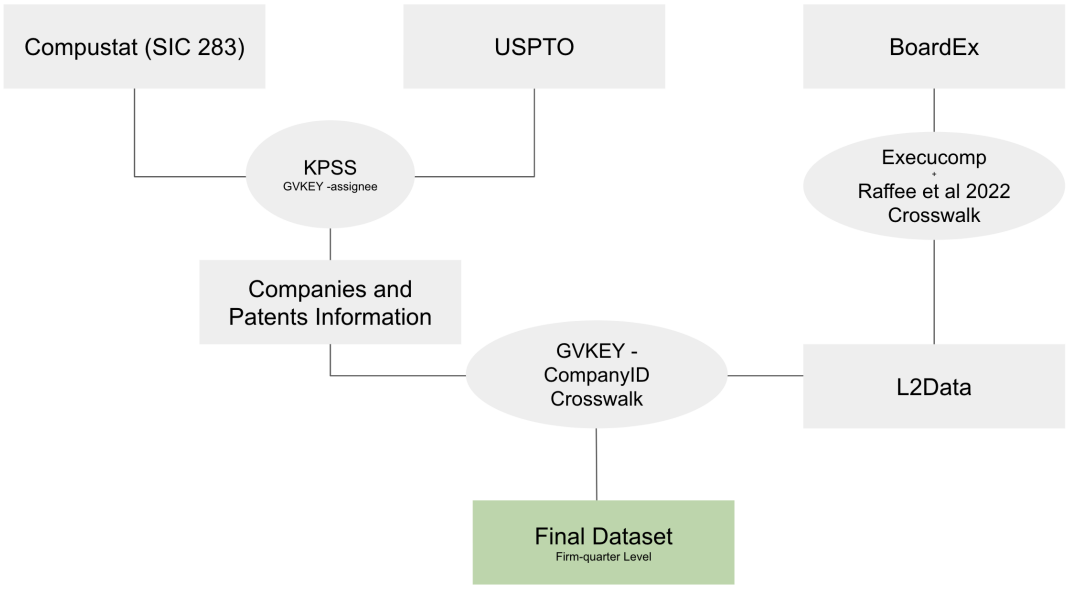

Figure 4: Dataset Construction Process.

This integration process in Figure 4 results in a panel dataset covering 5545 firm-period observations across 148 unique firms, providing a comprehensive foundation for analyzing the relationship between top management characteristics and female-focused innovation.

A2 Appendix: Face validity of the measure of female-focused innovation

To provide face validity to our measure of female-focused innovation, Tables 4 and 5 present a selection of patents labeled as female-focused and male-focused. For each patent, we report the title, abstract, and the top five Medical Subject Headings (MeSH) terms identified using the *MeSH on Demand tool*[4], developed by the U.S. National Library of Medicine. MeSH terms are standardized keywords used for indexing and retrieving biomedical information in databases such as PubMed. Each term belongs to a hierarchical structure known as the MeSH tree, which organizes concepts by specificity and thematic

[4] https://meshb.nlm.nih.gov/MeSHonDemand

area. By comparing the content of each patent to the MeSH terms most frequently associated with it, we offer an external point of reference that can help assess whether the thematic focus aligns with the assigned categorization.

Patent	Title	Abstract	MeSh Terms
10391105	Methods of treating certain depressive disorders and delirium tremens	The disclosure provides a method of treating a patient having *postpartum depression*, *premenstrual dysphoric disorder*, *menopausal depression*, or delirium tremens comprising administering an effective amount of an injectable neurosteroid formulation to the patient. In certain embodiments the injectable neurosteroid formulation is an injectable ganaxolone formulation. Two types of injectable neurosteroid formulations may be used in the disclosed methods. The first such formulation is an injectable formulation containing ganaxolone and sulfobutyl ether β-cyclodextrin in a 1:1 inclusion complex. The second such formulation is an injectable neurossteroid formulation comprising a neurosteroid, preferably ganaxolone, but which may be a neurosteroid selected from allopregnanolone, ganaxolone, alphaxalone, alphadolone, hydroxydione, minaxolone, pregnanolone, tetrahydrocorticosterone, and combinations of the foregoing; a surface stabilizer selected from hydroxyl ethyl starch, dextran, and povidone; and a surfactant.	Female, minaxolone, ganaxolone, Pregnanolone, hydroxydione
9006258	Method of treating *female sexual dysfunction* with a PDE1 inhibitor	The present invention relates to a new use for compounds that inhibit phosphodiesterase 1 (PDE1), e.g., that inhibit PDE1-mediated suppression of the dopamine D1 receptor and/or progesterone signaling pathways, including, e.g., methods of treatment or prophylaxis for conditions which may be ameliorated by enhancing the progesterone signaling response, particularly *female sexual dysfunction*.	Female, Phosphoric, Diester Hydrolases, Progesterone, Inventions, Signal Transduction
9358299	17-hydroxyprogesterone ester-containing oral compositions and related methods	The present invention provides for bioavailable oral dosage forms containing esters of 17-hydroxyprogesterone as well as related methods The oral dosage forms can be formulated for *pregnancy* support and can include a therapeutically effective amount of an ester of 17-hydroxyprogesterone and a pharmaceutically acceptable carrier In another embodiment a pharmaceutically acceptable oral dosage form for *pregnancy* support is provided The pharmaceutically acceptable oral dosage can include a therapeutically effective amount of an ester of 17-hydroxyprogesterone and a pharmaceutically acceptable carrier The oral dosage form can when measured using a USP Type-II dissolution apparatus in 900 mL of deionized water with 0 5 w v of sodium lauryl sulfate at 50 RPM at 37 C release at least 20 wt of the dose of the ester of 17-hydroxyprogesterone after 60 minutes or in the alternative release at least 20 wt more after 60 minutes than an equivalently dosed oral dosage form without the carrier.	Pregnancy, Female, Sodium, Dodecyl Sulfate, Drug Carriers, Water

Table 4: Patents categorized as Female-focused

Patent	Title	Abstract	MeSh Terms
8987319	Selective androgen receptor	This invention provides compounds of Formula (I), (Ia), (Ib), (Ic), (II), (IIa) or (III) and or salts thereof, pharmaceutical compositions comprising a compound of Formula (I), (Ia), (Ib), (Ic), (II), (IIa) or (III) and a pharmaceutically acceptable excipient, methods of modulating the androgen receptor, methods of treating diseases beneficially treated by an androgen receptor modulator (e.g., sarcopenia, *prostate cancer*, contraception, type II diabetes related disorders or diseases, anemia, depression, and renal disease) and processes for making compounds of Formula (I), (Ia), (Ib), (Ic), (II), (IIa) or (III) and intermediates useful in the preparation of same.	Male, Diabetes Mellitus Type 2, Excipients, Salts, Inventions
9387216	Biomarkers for treatment of neoplastic disorders using androgen-targeted therapies	Described herein are methods and compositions for the treatment of *prostate* cancer in a subject in need thereof. The prostate cancer may be a castration resistant and an androgen receptor antagonist-resistant prostate cancer. The methods may comprise administering to the subject a CYP17-lyase inhibitor of Formula II.	Male, Androgen Receptor Antagonists, Steroid 17-alpha-Hydroxylase, Prostatic Neoplasms and Castration-Resistant, Orchiectomy
6313308	Methods for the preparation of biphenyl isoxazole sulfonamides	Methods for the preparation of biphenyl isoxazole sulfonamides and intermediates therof. The present invention also relates to the novel intermediates prepared by these methods. The biphenyl isoxazole sulfonamides prepared by the present methods are endothelin antagonists useful, inter alia, for the treatment of hypertension, congestive heart failure and *male erectile dysfunction*.	Male, Endothelin Receptor Antagonists, Sulfonamides, Antihypertensive Agents, Erectile Dysfunction

Table 5: Patents categorized as Male-focused

References

Aiken, L. S., West, S. G., and Reno, R. R. (1991). *Multiple regression: Testing and interpreting interactions*. sage.

Bose, D., Segui-Gomez, ScD, M., and Crandall, J. R. (2011). Vulnerability of female drivers involved in motor vehicle crashes: an analysis of us population at risk. *American Journal of Public Health*, 101(12):2368–2373.

Brielmaier, C. and Friesl, M. (2023). The attention-based view: Review and conceptual extension towards situated attention. *International Journal of Management Reviews*, 25(1):99–129.

Brower, J. and Nath, P. (2018). Antecedents of market orientation: marketing ceos, cmos, and top management team marketing experience. *Marketing Letters*, 29(4):405–419.

Calder-Wang, S. and Gompers, P. A. (2021). And the children shall lead: Gender diversity and performance in venture capital. *Journal of Financial Economics*, 142(1):1–22.

Chandy, R., Hopstaken, B., Narasimhan, O., and Prabhu, J. (2006). From invention to innovation: Conversion ability in product development. *Journal of Marketing Research*, 43(3):494–508.

Chandy, R. K. and Tellis, G. J. (2000). The incumbent's curse? incumbency, size, and radical product innovation. *Journal of marketing*, 64(3):1–17.

Cifor, M. and Garcia, P. (2019). Inscribing gender: A duoethnographic examination of gendered values and practices in fitness tracker design. *Proceedings of the 52nd Hawaii International Conference on System Science*.

Corwin, E. S., Loncarich, H., and Ridge, J. W. (2022). What's it like inside the hive?

managerial discretion drives tmt gender diversity of women-led firms. *Journal of Management*, 48(4):1003–1034.

Cronqvist, H. and Yu, F. (2017). Shaped by their daughters: Executives, female socialization, and corporate social responsibility. *Journal of Financial Economics*, 126(3):543–562.

Cyert, R. and March, J. (1963). *A Behavioral Theory of the Firm.* Englewood Cliffs, NJ: Prentice-Hall.

Dasgupta, A., Ha, L., Jonnalagadda, S., Schmeiser, S., and Youngerman, H. (2018). The daughter effect: do ceos with daughters hire more women to their board? *Applied Economics Letters*, 25(13):891–894.

Day, G. S. and Nedungadi, P. (1994). Managerial representations of competitive advantage. *Journal of Marketing*, 58(2):31–44.

Derks, B., Van Laar, C., and Ellemers, N. (2016). The queen bee phenomenon: Why women leaders distance themselves from junior women. *The Leadership Quarterly*, 27(3):456–469.

Derks, B., Van Laar, C., Ellemers, N., and Raghoe, G. (2015). Extending the queen bee effect: How Hindustani workers cope with disadvantage by distancing the self from the group. *Journal of Social Issues*, 71(3):476–496.

Devlin, J., Chang, M.-W., Lee, K., and Toutanova, K. (2018). Bert: Pre-training of deep bidirectional transformers for language understanding. arXiv preprint arXiv:1810.04805.

Dezsö, C. L. and Ross, D. G. (2012). Does female representation in top management improve firm performance? a panel data investigation. *Strategic Management Journal*, 33(9):1072–1089.

Duguid, M. (2011). Female tokens in high-prestige work groups: Catalysts or inhibitors of group diversification? *Organizational Behavior and Human Decision Processes*, 116(1):104–115.

Duguid, M. M., Loyd, D. L., and Tolbert, P. S. (2012). The impact of categorical status, numeric representation, and work group prestige on preference for demographically similar others: A value threat approach. *Organization Science*, 23(2):386–401.

D'Angelo, J. K., Dunn, L., and Valsesia, F. (2024). Is this for me? Differential responses to skin tone inclusivity initiatives by underrepresented consumers and represented consumers. *Journal of Marketing*, page 00222429241268634.

Einiö, E., Feng, J., and Jaravel, X. (2019). Social push and the direction of innovation. *Available at SSRN 3383703*.

Ellemers, N., Van den Heuvel, H., De Gilder, D., Maass, A., and Bonvini, A. (2004). The underrepresentation of women in science: Differential commitment or the queen bee syndrome? *British Journal of Social Psychology*, 43(3):315–338.

Feng, H., Morgan, N. A., and Rego, L. L. (2015). Marketing department power and firm performance. *Journal of Marketing*, 79(5):1–20.

Finkelstein, S. (1992). Power in top management teams: Dimensions, measurement, and validation. *Academy of Management Journal*, 35(3):505–538.

Fordham, S. (2008). Beyond capital high: On dual citizenship and the strange career of "acting white". *Anthropology & Education Quarterly*, 39(3):227–246.

Gagliardi, L. and Novelli, E. (2024). Female entrepreneurs targeting women: Strategic redirection under scientific decision-making. *Organization Science*.

Glynn, A. N. and Sen, M. (2015). Identifying judicial empathy: Does having daughters cause judges to rule for women's issues? *American Journal of Political Science*, 59(1):37–54.

Gorman, E. H. (2005). Gender stereotypes, same-gender preferences, and organizational variation in the hiring of women: Evidence from law firms. *American Sociological Review*, 70(4):702–728.

Gould, J. A., Kulik, C. T., and Sardeshmukh, S. R. (2018). Trickle-down effect: The impact of female board members on executive gender diversity. *Human Resource Management*, 57(4):931–945.

Graves Jr, J. L., Kearney, M., Barabino, G., and Malcom, S. (2022). Inequality in science and the case for a new agenda. *Proceedings of the National Academy of Sciences*, 119(10):e2117831119.

Guldiken, O., Mallon, M. R., Fainshmidt, S., Judge, W. Q., and Clark, C. E. (2019). Beyond tokenism: How strategic leaders influence more meaningful gender diversity on boards of directors. *Strategic Management Journal*, 40(12):2024–2046.

Gupta, H. (2022). Do female researchers increase female enrollment in clinical trials? Technical report, Working paper.

Hambrick, D. C. (2007). Upper echelons theory: An update. *Academy of management review*, 32(2):334–343.

Hambrick, D. C. and Mason, P. A. (1984). Upper echelons: The organization as a reflection of its top managers. *Academy of management review*, 9(2):193–206.

Harmancioglu, N., Grinstein, A., and Goldman, A. (2010). Innovation and performance outcomes of market information collection efforts: The role of top management team involvement. *International Journal of Research in Marketing*, 27(1):33–43.

Heckman, J. J. (1976). The common structure of statistical models of truncation, sample selection and limited dependent variables and a simple estimator for such models. In *Annals of economic and social measurement, volume 5, number 4*, pages 475–492. NBER.

Heckman, J. J. (1979). Sample selection bias as a specification error. *Econometrica: Journal of the econometric society*, pages 153–161.

Iyer, J. K. (2020). From the boardroom to the consulting room, we need women in pharma. *World Economic Forum*.

Johnson, P. O. and Neyman, J. (1936). Tests of certain linear hypotheses and their application to some educational problems. *Statistical Research Memoirs*.

Kabátek, J. and Ribar, D. C. (2021). Daughters and divorce. *The Economic Journal*, 131(637):2144–2170.

Kemble, E., Pérez, L., Sartori, V., Tolub, G., and Zheng, A. (2022). Unlocking opportunities in women's healthcare. McKinsey & Company. Last accessed on 13/05/22.

Koning, R., Samila, S., and Ferguson, J.-P. (2020). Inventor gender and the direction of invention. In *AEA Papers and Proceedings*, volume 110, pages 250–254. American Economic Association, Nashville, TN 37203.

Koning, R., Samila, S., and Ferguson, J.-P. (2021). Who do we invent for? Patents by women focus more on women's health, but few women get to invent. *Science*, 372(6548):1345–1348.

Kozlowski, D., Larivière, V., Sugimoto, C. R., and Monroe-White, T. (2022). Intersectional inequalities in science. *Proceedings of the National Academy of Sciences*, 119(2):e2113067119.

Li, Q., Maggitti, P. G., Smith, K. G., Tesluk, P. E., and Katila, R. (2013). Top management attention to innovation: The role of search selection and intensity in new product introductions. *Academy of Management Journal*, 56(3):893–916.

Lyngsie, J. and Foss, N. J. (2017). The more, the merrier? women in top-management teams and entrepreneurship in established firms. *Strategic management journal*, 38(3):487–505.

Mickey, E. L. and Smith-Doerr, L. (2022). Gender and innovation through an intersectional lens: Re-imagining academic entrepreneurship in the united states. *Sociology Compass*, 16(3):e12964.

Nielsen, M. W., Bloch, C. W., and Schiebinger, L. (2018). Making gender diversity work for scientific discovery and innovation. *Nature Human Behaviour*, 2(10):726–734.

Ocasio, W. (1997). Towards an attention-based view of the firm. *Strategic Management Journal*, 18(S1):187–206.

Ocasio, W. (2011). Attention to attention. *Organization Science*, 22(5):1286–1296.

Oswald, A. J. and Powdthavee, N. (2010). Daughters and left-wing voting. *The Review of Economics and Statistics*, 92(2):213–227.

Petrin, A. and Train, K. (2010). A control function approach to endogeneity in consumer choice models. *Journal of marketing research*, 47(1):3–13.

Post, C. and Byron, K. (2015). Women on boards and firm financial performance: A meta-analysis. *Academy of Management Journal*, 58(5):1546–1571.

Prabhu, J. C., Chandy, R. K., and Ellis, M. E. (2005). The impact of acquisitions on innovation: Poison pill, placebo, or tonic? *Journal of Marketing*, 69(1):114–130.

Radford, A., Narasimhan, K., Salimans, T., and Sutskever, I. (2018). Improving language understanding by generative pre-training. Technical Report, OpenAI.

Raffiee, J., Fehder, D., and Teodoridis, F. (2022). Revealing the revealed preferences of public firm ceos and top executives: A new database from credit card spending. *Strategic Management Journal*, 43(10):2042–2065.

Raley, S. and Bianchi, S. (2006). Sons, daughters, and family processes: Does gender of children matter? *Annual Review of Sociology*, 32(1):401–421.

Sands, S. and Ferraro, C. (2025). DEI rollbacks: Consequences and considerations for brand managers. *Business Horizons*.

Sorescu, A. B., Chandy, R. K., and Prabhu, J. C. (2003). Sources and financial consequences of radical innovation: Insights from pharmaceuticals. *Journal of marketing*, 67(4):82–102.

Spencer, R. J., Rice, L. W., Ye, C., Woo, K., and Uppal, S. (2019). Disparities in the allocation of research funding to gynecologic cancers by funding to lethality scores. *Gynecologic Oncology*, 152(1):106–111.

Srivastava, C., Kashmiri, S., and Mahajan, V. (2023). Customer orientation and financial performance: Women in top management teams matter! *Journal of Marketing*, 87(2):190–209.

Staines, G., Tavris, C., and Jayaratne, T. E. (1974). The queen bee syndrome. *Psychology Today*, (7):55–60.

Stroube, B. K. and Waguespack, D. M. (2024). Status and consensus: Heterogeneity in audience evaluations of female-versus male-lead films. *Strategic Management Journal*, 45(5):994–1024.

Talke, K., Salomo, S., and Rost, K. (2010). How top management team diversity affects innovativeness and performance via the strategic choice to focus on innovation fields. *Research Policy*, 39(7):907–918.

Torchia, M., Calabrò, A., and Huse, M. (2011). Women directors on corporate boards: From tokenism to critical mass. *Journal of Business Ethics*, 102:299–317.

Truffa, F. and Wong, A. (2024). Undergraduate gender diversity and the direction of scientific research.

Umashankar, N., Bahadir, S. C., and Bharadwaj, S. (2022). Despite efficiencies, mergers and acquisitions reduce firm value by hurting customer satisfaction. *Journal of Marketing*, 86(2):66–86.

Varma, R., Bommaraju, R., and Singh, S. S. (2023). Female chief marketing officers: When and why do their marketing decisions differ from their male counterparts'? *Journal of Marketing Research*, 60(6):1154–1176.

Warner, R. L. and Steel, B. S. (1999). Child rearing as a mechanism for social change: The relationship of child gender to parents' commitment to gender equity. *Gender & Society*, 13(4):503–517.

Washington, E. L. (2008). Female socialization: How daughters affect their legislator fathers' voting on women's issues. *American Economic Review*, 98(1):311–332.

Wilcox, A. J., Weinberg, C. R., and Baird, D. D. (1995). Timing of sexual intercourse in relation to ovulation—effects on the probability of conception, survival of the pregnancy, and sex of the baby. *New England Journal of Medicine*, 333(23):1517–1521.

Wu, Z., Naldi, L., Wennberg, K., and Uman, T. (2024). Learning from their daughters: family exposure to gender disparity and female representation in male-led ventures. *Management Science*, 70(2):671–693.

Yadav, M. S., Prabhu, J. C., and Chandy, R. K. (2007). Managing the future: Ceo attention and innovation outcomes. *Journal of Marketing*, 71(4):84–101.

You, Y., Srinivasan, S., Pauwels, K., and Joshi, A. (2020). How CEO/CMO characteristics affect innovation and stock returns: findings and future directions. *Journal of the Academy of Marketing Science*, 48:1229–1253.

Zhou, S., Chai, S., and Freeman, R. B. (2024). Gender homophily: In-group citation preferences and the gender disadvantage. *Research Policy*, 53(1):104895.

Chapter 2

A Day in the Life: Mapping Brand DNA and Local Identities Through Multimodal User-Generated Content

Margherita Caprara[1],**Francesco Grossetti**[2], **Paola Cillo**[3], **Gaia Rubera**[1]

[1]Department of Marketing, Bocconi University

[2]Department of Accounting and BIDSA, Bocconi University

[3]Department of Management and Technology, Bocconi University

This study introduces Multimodal Brand Community Analysis (MBCA), a scalable framework that integrates computer vision and network analysis to capture how brands are embedded in consumers' everyday lives. Using more than 200 million Instagram images from eight global cities, we detect brand logos, identify co-occurring hashtags in captions, and generate interpretable descriptions of brand communities with generative AI. Applying MBCA to Adidas and Nike, we document three insights. First, by integrating visual and textual modalities, MBCA uncovers richer and more interpretable consumer communities than either modality alone. Second, brand meaning is not monolithic but varies systematically across cities, with some themes recurring globally as part of a brand's DNA and others reflecting local lifestyles and cultural contexts. Third, we reveal context bias: consumers disproportionately mention brands in aligned contexts (e.g., sports for athletic brands) while overlooking routine uses where brands are present but not explicitly tagged. Theoretically, our work advances multimodal analysis in marketing, introduces a scalable approach to localized brand meaning, and extends UGC research beyond declarative expressions to incidental brand presence. Managerially, MBCA equips marketers with a data-driven toolkit to detect unprompted brand use, localize campaigns at the city level, and translate UGC-based insights into context-aware creative execution.

Keywords: Customer base analysis, social media, user-generated content, artificial intelligence, computer vision, community detection

2.1 Introduction

Understanding how consumers use brands in everyday life remains a central objective in marketing research. User-generated content (UGC) offers an unprecedented opportunity to do so. Not surprisingly, UGC has become a prominent focus in the marketing literature and firms invest heavily in analyzing it (Schwarz, 2023). However, despite this investment, managers increasingly express disappointment with the insights delivered. They complain that conventional UGC analyses produce fragmented indicators (e.g., like counts or sentiment), leaving them with insights too shallow to inform decision making (Hollow, 2023; Kushner, 2025). Reflecting these frustrations, both practitioners and academics have urged a shift toward more interpretive, context-aware approaches, urging researchers to "think like anthropologists" by treating social media as a cultural field site where brand meaning emerges from immersion in the richness of UGC (Fournier et al., 2016). However, current marketing research has not yet equipped managers with the methodological tools to realize this promise, leaving UGC analyses short of their potential.

Specifically, existing UGC research suffers from three important gaps. First, most studies identify brand-related content through explicit references such as hashtags, keywords, or captions (Nam et al., 2017; Klostermann et al., 2018; Liu et al., 2020). This approach captures only moments when consumers *deliberately* foreground the brand, overlooking unprompted uses in which brands appear incidentally in everyday life. Second, prior work generally examines a single modality—text or images—even though consumer expression online is inherently multimodal. Text reveals what consumers choose to articulate, while images capture the lifestyle contexts in which brands are embedded; isolating one modality risks omitting critical signals. Third, most analyses collapse UGC across geographies to produce a global brand understanding (Lee and Bradlow, 2011; Netzer et al., 2012; Yang et al., 2022), erasing the cultural variation in brand meaning across

places. Together, these limitations highlight the need for a new methodology that can capture unprompted brand use in consumers' everyday life and recover both global and local dimensions of brand use.

We address this gap by introducing a scalable multimodal framework that combines computer vision, Natural Language Processing (NLP), and community detection to identify and interpret brand related clusters of geo-located consumers posting UGC on Instagram. We term our approach *Multimodal Brand Community Analysis* (MBCA), which consists of four steps. First, we collect Instagram images that are geo-tagged to specific cities. Second, we detect brand presence in these images using visual logo recognition. Third, we identify consumer communities by applying community-detection algorithms to co-occurrence networks of hashtags in images containing a given brand's logo. Fourth, we generate interpretable labels and descriptions for each community using a large language model. Together, these steps yield consumer communities that reflect shared preferences, interests, and lifestyles. MBCA is automated, scalable, and produces human-interpretable outputs, enabling researchers to discover consumer communities across geographies. To illustrate its value, we apply MBCA to a dataset of over 200 million images collected across eight global cities, focusing on two athletic brands, Adidas and Nike. We analyze consumer communities at the city level and compare results under two sampling regimes: logo-based sampling (where images are included based on visual brand presence, regardless of textual mention) and mention-based sampling (where images are included based on explicit textual references to the brand).

Our analysis yields three main findings. First, by integrating visual and textual modalities, we uncover richer and more interpretable communities of brand use than either modality alone could reveal. Images capture the unspoken, behavioral contexts in which brands are embedded, while hashtags provide the semantic cues that situate those contexts. Together, they offer a holistic view of how consumers use and experience brands. Second, brand use is not monolithic across cities. Prior research in global branding has

contrasted standardized and localized positioning across countries (e.g., Alden et al. 1999; Steenkamp 2001), while most UGC studies collapse consumer content into a single global representation of brand associations (e.g., Hartmann et al. 2021; Klostermann et al. 2018; Rust et al. 2021). Both perspectives overlook the finer-grained, city-specific patterns of brand use within countries, as consumers integrate brands into everyday routines differently across locations. We show that while certain themes recur globally—what we term a "brand's DNA"—brand associations also differ systematically across cities, even between geographically proximate markets such as San Francisco and Los Angeles. By documenting these systematic city-level differences, our study provides a scalable framework to complement prior qualitative research (e.g., Thompson and Arsel 2004) that, while culturally rich, is costly and difficult to replicate across multiple locations. Third, we find that relying exclusively on brand mentions captures only moments in which consumers explicitly position a brand at the center of their message, while overlooking how brands are woven into everyday life without being directly highlighted. We call this omission "context bias". Such bias produces an incomplete, and at times distorted, view of consumer–brand interactions.

Managerially, our method offers brand managers a scalable, data-driven approach to cultural localization. By moving beyond explicit brand mentions to capture unprompted visual brand use, it uncovers use contexts and lifestyle communities that traditional UGC analysis overlook. Our city-level framework enables marketers to detect both global brand themes and locally distinctive expressions, supporting the design of marketing campaigns that resonate more deeply with local audiences. Finally, by structuring consumer insights into semantically coherent, model-ready prompts, our approach bridges the gap between insight generation and creative execution. We show how managers can feed these outputs into generative AI systems to design city-specific campaigns, directly shaping and adjusting their tone, imagery, and messaging, while reducing friction between research, strategy, and creative functions. In this way, our method transforms UGC into a strategic asset for

context-aware brand building in an increasingly fragmented and dynamic marketplace.

2.2 Literature on User Generated Content

The significance of UGC in providing customer insights, particularly for understanding customer needs, preferences, opinions, and motivations, is well-documented in recent marketing research (Chung et al., 2022; Culotta and Cutler, 2016; Dang et al., 2025; Hartmann et al., 2021; Klostermann et al., 2018; Liu et al., 2020; Moe et al., 2017; Nam et al., 2017; Nam and Kannan, 2014; Rust et al., 2021; Timoshenko and Hauser, 2019). In Table 2.1, we provide an overview of previous articles in this area.

This body of work has primarily analyzed one type of data at a time—either numeric, textual, or visual content. For instance, past research has used numeric data such as the number of likes or comments (Yang et al., 2022) or Twitter followership (Culotta and Cutler, 2016) to retrieve brand perceptions and market structure. Previous studies have also extensively analyzed textual data, including product reviews and social media posts, to track brand reputation (Rust et al., 2021), uncover consumer needs and trends (Timoshenko and Hauser, 2019), measure sales conversion (Liu et al., 2019), develop segmentation techniques (Jagabathula et al., 2018; Tirunillai and Tellis, 2014) as well as to create market structure representations (Lee and Bradlow, 2011; Netzer et al., 2012). Similarly, researchers have used social tags to categorize and contextualize consumer-posted content, enabling the measurement of brand equity and perceptions (Nam and Kannan, 2014; Nam et al., 2017). Just recently, the marketing literature has started analyzing visual UGC content. For instance, Zhang and Luo (2023) estimate restaurants' survival chances from customer-posted images on Yelp. Klostermann et al. (2018) analyze the text of Instagram images containing the hashtag #mcdonalds. Similarly, Liu et al. (2020) map images hashtagged with brand names to pre-defined attributes (i.e., glamorous, rugged, healthy, fun) to measure how brands are portrayed along these attributes.

Paper	Type of Data	Source of Data	Geographic scope	Marketing Concept
Culotta and Cutler (2016)	Numeric	Twitter followership	Global	Brand perceptions
Hartmann et al. (2021)	Visual and Text	Branded images, captions, and comments	Global	Brand engagement and purchase intention
Jagabathula et al. (2018)	Numeric	Clicks, purchase, demographics, and customers' metrics	N.A.	Segmentation
Klostermann et al. (2018)	Visual and Text	Social tags and Captions from images posted	Global	Brand perceptions
Lee and Bradlow (2011)	Text	Product review	Global	Market structure
Liu et al. (2020)	Visual	Instagram images	Global	Brand perceptions
Liu et al. (2019)	Text	Product review	Country	Sales conversions
Nam and Kannan (2014)	Text	Social tags	Country	Brand equity
Nam et al. (2017)	Text	Social tags	Global	Brand perceptions
Netzer et al. (2012)	Text	Web forums	Country	Market structure
Rust et al. (2021)	Text	Twitter comments	Global	Brand reputation
Timoshenko and Hauser (2019)	Text	Product review	N.A.	Customers' needs
Tirunillai and Tellis (2014)	Text	Product review	N.A.	Segmentation
Wang et al. (2022)	Text	Product review	N.A.	Market structure
Yang et al. (2022)	Numeric	Likes and comments	Global	Market structure
Zhang and Luo (2023)	Visual	Restaurant images on Yelp	Country	Firm survival
This study	Visual and Text	Branded images and social tags	City	Habits, preferences, and lifestyle of a brand's customers

Note: N.A. stands for not applicable.

Most of this literature shares three key limitations. First, nearly all studies rely on explicit brand references to identify relevant UGC, typically through hashtags, keywords, or captions (Klostermann et al., 2018; Liu et al., 2020; Nam et al., 2017). While effective for identifying intentional brand talk, this approach systematically privileges moments when consumers deliberately foreground the brand. It risks overlooking more routine, unprompted uses in which brands appear incidentally in everyday life, but are not actively mentioned. As a result, current research may produce a partial picture of how consumers engage with brands. In contrast, our multimodal approach uses images to identify brand presence and hashtags to interpret usage patterns. In this way, we can capture instances of unprompted brand use, including moments where consumers use the brand into their everyday lives without intending to draw attention to it. By capturing these less intentional, more naturalistic interactions, our approach offers a broader, and potentially more representative, view of brand use. Differently, mention-based sampling, whether applied to textual or visual UGC, systematically overlooks these interactions, as our analysis shows.

Second, most prior work focuses on a single type of UGC, often analyzing either text or images in isolation. However, consumer expression on social media is inherently multimodal, and focusing on a single modality risks omitting key signals. Textual data, such as hashtags, captures what consumers choose to articulate. Differently, visual content reflects behaviors and lifestyle contexts that are often implicit and not deliberately communicated. By integrating both visual and textual modalities, our approach captures a more complete representation of who a brand's consumers are, based on the contexts in which they use the brand in everyday life. Visual UGC allows us to detect brand presence even when consumers do not mention it, while the accompanying hashtags help contextualize that presence in terms of interests and habits. This multimodal integration enables us to identify recurring patterns in who uses a brand and in what contexts. Hartmann et al. (2021) is closely related to our study. The authors use a pre-trained convolutional

neural network to classify selfies with brand logo into three categories: consumer selfies, brand selfies, and packshots. They report that consumer-selfie images generate more likes and comments, while brand selfies associate with higher purchase intentions, which the authors retrieve from comments to the posted picture. Similarly, our study also uses both visual and text data. However, we analyze all images containing brand logos, not just selfies in which the branded product is clearly visible. Since brands frequently encourage consumers to take selfies using the branded product (Hartmann et al., 2021), our approach allows us to observe consumers using brands in a more natural setting. Further, we are interested in using images to understand the lifestyle and habits of a brand's consumer base, while Hartmann et al. (2021) study how different types of selfies generate engagement.

Third, existing studies typically analyze UGC in the aggregate, across geographies. For example, prior studies derive brand relationships by mining textual UGC and aggregating co-occurrence patterns into perceptual or competitive maps (Netzer et al., 2012; Lee and Bradlow, 2011; Yang et al., 2022). While these methods offer valuable tools for understanding competitive proximity, they do so by collapsing all UGC into a single, undifferentiated brand space, overlooking how brand associations may vary across geographies. Differently, our approach analyzes brand-related UGC at the city level. In this way, we can identify both global themes, which we refer to as the "brand DNA", and local ones, which reflect how consumers in different cities use and associate with the brand. These insights are essential for context-aware marketing strategy. Summarizing, the limitations in past studies bias marketers' understanding of a brand's customer base. Mention-based sampling excludes consumers who engage with the brand without self-identifying. Single-modality analyses risk missing either the presence of the brand (if relying only on text) or the context of brand use (if relying only on images). Global aggregation ignores the different ways in which brands are used across places. Our multimodal, city-level approach addresses these gaps and produces richer, more actionable insights about who a brand's

customers are and how they use the brand in their everyday lives.

2.3 Data and Methods

MBCA comprises four steps, which we illustrate in Figure 2.1 and now describe in detail.

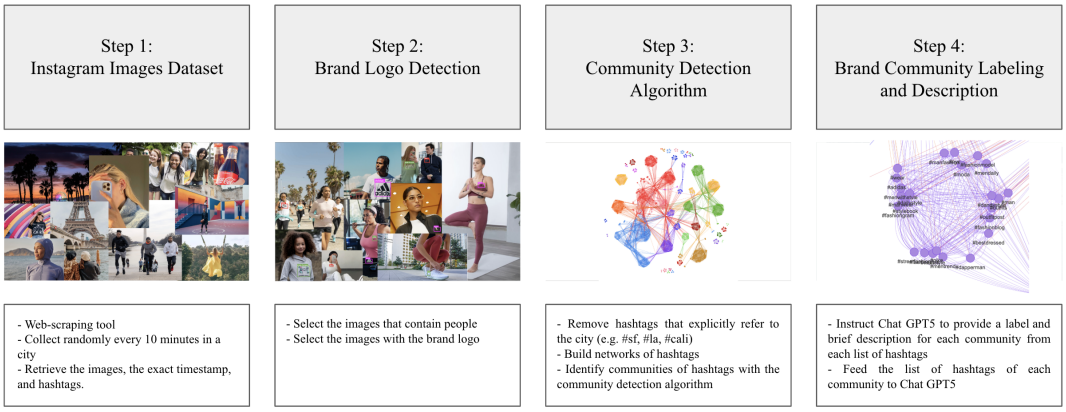

Figure 2.1: Illustration of the Four Multimodal Brand Community Analysis Steps.

2.3.1 Instagram Images Dataset

The first step in our analysis is the collection of user-generated images. For this purpose, we developed a web-scraping tool that systematically gathers Instagram images and their metadata. Our dataset spans 2015-2018 and covers eight major U.S. and European cities: Atlanta, Chicago, Cleveland, London, Los Angeles, Miami, Paris, and San Francisco. We focus on this period for three reasons. First, it provides a large and stable window of consumer activity before the rapid platform diversification that began in 2019 with the rise of TikTok and the shift toward short-form video (e.g., reels). By restricting analysis to a period where Instagram was the dominant global image-sharing platform, we minimize confounds arising from transformations in social media usage patterns and ensure comparability across cities and brands. Second, the 2015–2018 window precedes the COVID-19

pandemic, which profoundly disrupted consumer routines and social media behavior. Our design isolates brand-related communities in a relatively stable cultural and market environment, avoiding distortions caused by an unprecedented global shock. Third, the size and multimodality of the dataset (over 200 million images across eight cities) required significant computational resources for image recognition, network construction, and community detection. Focusing on this time frame struck a balance between breadth, depth, and feasibility of large-scale multimodal analysis.

Our scraping tool samples batches of 100 images every 10 minutes in each city and collects metadata including, timestamps, captions, and hashtags. Importantly, the scraping mechanism is designed to capture images regardless of whether they contain explicit brand references (through hashtags or mentions). This strategy enables us to observe brand use in more natural contexts, rather than restricting analysis to posts where consumers intentionally highlight a brand. The resulting dataset, more than 200 million images across eight cities, provides a uniquely rich base for studying brand use at scale.

2.3.2 Brand Logo Detection

To identify images containing brand logos, we implement a two-step approach. First, we retain only images that contain at least one person. This step filters out pictures of storefronts, billboards, or product displays where logos may be visible but do not reflect actual consumer use. Restricting the sample to images with people increases confidence that we are analyzing UGC portraying active brand use rather than passive exposure. We detect people using YOLO (You Only Look Once) v5 (Redmon and Farhadi, 2018; Ultralytics, 2021), a state-of-the-art deep neural network architecture trained on large-scale datasets of images such as ImageNet (Deng et al., 2009). Its unified, end-to-end design allows for real-time object localization and classification in a single pass, offering high accuracy with low false positives.

Second, we fine-tune YOLOv5 on a custom dataset of diverse brand logos, consist-

ing of 13,500 manually annotated Instagram images, with a minimum of 500 annotated examples for each logo type (e.g., Adidas has three primary logo variants, yielding approximately 1,500 images). We use the VGG Image Annotator (Dutta and Zisserman, 2019), a web-based platform for efficient manual labeling. One co-author and two research assistants independently annotate all images, and a second manual review resolves any conflicts to finalize the dataset. We compare two YOLOv5 variants—YOLOv5l and YOLOv5x—and opt for YOLOv5x as it delivers superior precision in identifying small, complex, or partially obscured logos, despite requiring more computational resources. Figure 2.2 illustrate the model's performance in detecting Adidas logos. As shown, YOLOv5x accurately identifies logos across varied conditions, including mirrored orientations and cluttered visual scenes, underscoring its robustness for large-scale brand logo recognition.

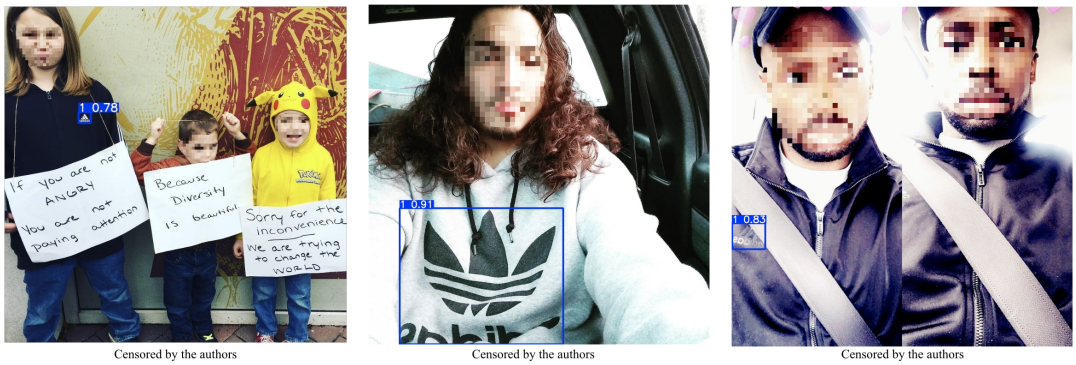

Censored by the authors Censored by the authors Censored by the authors

Figure 2.2: Prediction of Adidas Logos Using the Fine-Tuned YOLOv5x Architecture.

Like other object detection tasks, logo detection presents several challenges. Logos may be occluded (partially blocked by other objects), deformed (appearing in different orientations or perspectives), or subject to scale variation (appearing larger or smaller depending on distance from the camera). In addition, illumination changes, such as shadows, glare, or uneven lighting, can reduce visibility and hinder accurate detection. A further challenge lies in choosing the optimal confidence threshold, λ, which determines whether a detected object is classified as a logo. In YOLOv5x, confidence represents the model's certainty that a bounding box contains a specific class of object. It is computed

as the product of two components: the *objectness score* (likelihood that the bounding box contains an object at all) and the *class probability* (likelihood that the object belongs to a particular logo class).[1] Setting λ too low increases the risk of false positives, where non-logo objects are misclassified as logos. Setting it too high increases false negatives, where genuine logos go undetected. To balance there trade-offs, we calibrate λ using the F1 score, which captures the harmonic mean of precision and recall. We select the threshold that maximizes the F1 score in our application, though researchers as well as managers may adjust this value depending on their specific context (e.g., prioritizing recall to capture as many logos as possible versus precision to minimize misclassifications).

2.3.3 Community Detection Algorithm

With a clean dataset of images featuring various brand logos, the next step is to identify distinct groups of brand customers. Following previous research (e.g., Klostermann et al., 2018), we focus on hashtags, which capture the essence of online textual content (Nam et al., 2017). Instagram captions are typically brief and dominated by hashtags, making them an efficient data source for the community detection algorithm without sacrificing informational richness. To refine the corpus, we remove city-specific hashtags (e.g., "#cali", "#la", "#sf") which provide redundant information since location is already encoded in the metadata and add little to understanding brand-related content or consumer behavior. The resulting hashtag corpus forms the basis for constructing brand-associative networks at the city level.

We represent these associations using a *feature co-occurrence matrix* (FCM), where entries capture whether two hashtags appear together in the same caption of a given image. The FCM defines a graph structure $G = (V, E)$, with vertices $V = \{v_1, \ldots, v_n\}$ corresponding to hashtags and edges $E = \{(v_i, v_j) \mid v_i, v_j \in V, v_i \neq v_j\}$ representing hash-

[1] In YOLOv5, the confidence score ranges from 0 to 1, where higher values indicate greater certainty in a prediction. Specifically, λ = Objectness Score $\times$ Class Probability.

tag co-occurrences. Because co-occurrence has no inherent direction, the resulting graph G is undirected. To detect communities within G, we employ the walktrap algorithm (Pons and Latapy, 2005), a hierarchical method that has proven effective for extracting explainable and meaningful network communities (Golino and Epskamp, 2017). The algorithm is based on the intuition that random walks of short length are more likely to remain within the same community than to cross into others. A transition probability matrix P governs these walks, such that:

$$P_{ij} = \frac{A_{ij}}{d_i}, \tag{2.1}$$

with A denoting the adjacency matrix and d_i is the degree of vertex v_i which informs us about how many edges are connected to that vertex.[2] The probability distribution of a random walk starting from vertex u after l steps is denoted $p_u^{(l)}$. The distance between two vertices u and v is then defined as the Euclidean distance between their distributions:

$$d(u, v) = \left\| p_u^{(l)} - p_v^{(l)} \right\|. \tag{2.2}$$

Initially, each vertex forms its own community. The algorithm then iteratively merges communities C_u and C_v that minimize the distance:

$$d(C_u, C_v) = \frac{1}{|C_u| + |C_v|} \left[\sum_{i \in C_u,\, j \in C_v} d(i, j) \right], \tag{2.3}$$

where $|C_u|$ and $|C_v|$ denote the sizes of the two communities. The merging continues until all vertices form a single community, yielding a hierarchical dendrogram. The final partition is chosen to maximize modularity, defined as:

[2] The adjacency matrix of a graph is a square matrix used to represent the connections between vertices (i.e., hashtags) in the graph. If the graph has n vertices, then A is a $n \times n$ matrix where each element A_{ij} is defined as:

$$A_{ij} = \begin{cases} 1 & \text{if there is an edge between vertex } v_i \text{ and } v_j, \\ 0 & \text{otherwise.} \end{cases}$$

$$Q = \frac{1}{2m} \sum_{i,j} \left[A_{ij} - \frac{k_i k_j}{2m} \right] \delta(C_i, C_j), \tag{2.4}$$

where k_i and k_j are the degrees of vertices i and j, m is the total number of edges, and $\delta(C_i, C_j) = 1$ if vertices i and j belong to the same community and 0 otherwise. Maximizing modularity is effective in this context because it ensures that the detected communities are not merely artifacts of local noise but represent statistically meaningful clusters of consumer behavior. A high modularity score indicates that hashtags within a community co-occur far more frequently than would be expected under random chance, while connections across communities are relatively sparse. This makes the resulting partition especially well-suited for uncovering thematic cohesion within consumer groups: hashtags about fitness, fashion, or nightlife cluster together in meaningful ways, while weak links across themes are minimized. In this way, modularity provides an interpretable criterion that balances the algorithm's flexibility with empirical rigor, ensuring that the identified communities reflect genuine cultural and behavioral structures.

In practice, we run the walktrap algorithm iteratively across values of the walk length l. Smaller l yields fine-grained communities, while larger l produces coarser structures. By conducting a grid search, we identify the specification that maximizes modularity in Equation (2.4), ensuring the most cohesive community structures. This process generates sets of hashtags with strong internal linkages and weak external ties, thereby producing distinct thematic groups. The result is a detailed mapping of brand communities across cities, each characterized by a unique configuration of hashtags that captures localized consumer themes and interests.

2.3.4 Brand Community Labeling and Description

After identifying communities of hashtags that are densely interconnected, the next step is to label and interpret each community. We accomplish this using ChatGPT-5, an

advanced GenAI model capable of interpreting the colloquial, multilingual, and culturally embedded language common on Instagram. This step addresses a persistent limitation in both research and practice: when analyzing UGC across multiple cities, scholars and managers rarely possess the local cultural knowledge required to decode city-specific slang, expressions, or local references. Unlike human coders who may be unfamiliar with the slang of Cleveland or the lifestyle cues of Paris, GenAI can draw on a vast corpus of text and cultural knowledge to provide contextually grounded interpretations. In this sense, GenAI functions as a form of scalable digital ethnography. Traditional ethnographic fieldwork can uncover nuanced, culture-specific meanings, but it is costly, time-intensive, and difficult to replicate across multiple locations. By contrast, ChatGPT-5 approximates the interpretative capacity of ethnographers while applying it simultaneously across many cities and communities.

This capability is especially valuable given the scale and complexity of UGC. Many communities contain hundreds of hashtags that blend multiple languages, niche references, and shifting cultural cues. Left unstructured, such data is nearly impossible to parse manually. ChatGPT-5 translates this heterogeneity into coherent summaries and culturally sensitive descriptions, distilling both broad themes (e.g, fitness, nightlife, activism) and subtle nuances (e.g., local artists, neighborhood identities, underground music collectives). For instance, in San Francisco, one Adidas community centered around hip-hop culture featured hashtags like #thepiratedjs (a local DJ collective) and #kmeldjs (a Bay Area urban contemporary radio crew). While these tags clearly signal music and *DJing*, their specific connection to hip-hop culture is not immediately evident to an outsider. A human analyst unfamiliar with San Francisco's culture would likely overlook this link, and manually researching each hashtag would be prohibitively time-consuming. In contrast, ChatGPT-5 effortlessly situates these hashtags within the broader hip-hop context, producing a narrative that highlights the community's cultural orientation. By automating the interpretive process through GenAI, our approach makes it possible to

extract culturally grounded, human-interpretable insights from UGC at scale.

2.4 Application of MBCA

To illustrate the potential of our methodology, we apply MBCA to two iconic athletic brands: Nike and Adidas. We select these brands for two reasons. First, both are global players, allowing us to retrieve a large and diverse set of images containing their logos from cities around the world. Second, they dominated the athletic sector during our study period, with Nike representing a leading American brand and Adidas a leading European counterpart. This selection is purely illustrative as marketers can readily apply MBCA to any brand of interest, regardless of sector and geography.

For this application, we define an image as containing a brand logo when the YOLOv5x detection confidence meets or exceeds $\lambda \geq 0.70$, a threshold that maximizes the F1 score in our calibration. Importantly, λ is an adjustable parameter, giving researchers and managers the flexibility to fine-tune the balance between precision and recall to suit specific analytical goals.

We showcase MBCA in two complementary analyses: (a) within-city, between-brands, by comparing Nike and Adidas communities in the same city; and (b) within-brand, between-cities, by examining how a single brand's communities vary across different cities. We further compare communities identified through visual brand logo detection with those identified using the simpler method of explicit brand mentions, thereby demonstrating the added value of incorporating visual information. Due to space limitations, we present results for the top 10 communities by image count, which together account for 61-76% of images in each city. This choice balances interpretability with comprehensiveness: while MBCA is capable of detecting smaller niche communities, such components are often difficult to label meaningfully. Researchers and managers, however, can adjust the

number of communities analyzed depending on their objectives.[3] Figure 2.3 illustrates an example network of Nike communities in Cleveland, highlighting the top 10 clusters. This result clearly shows MBCA's ability to reveal both global brand patterns and place-specific cultural expressions.

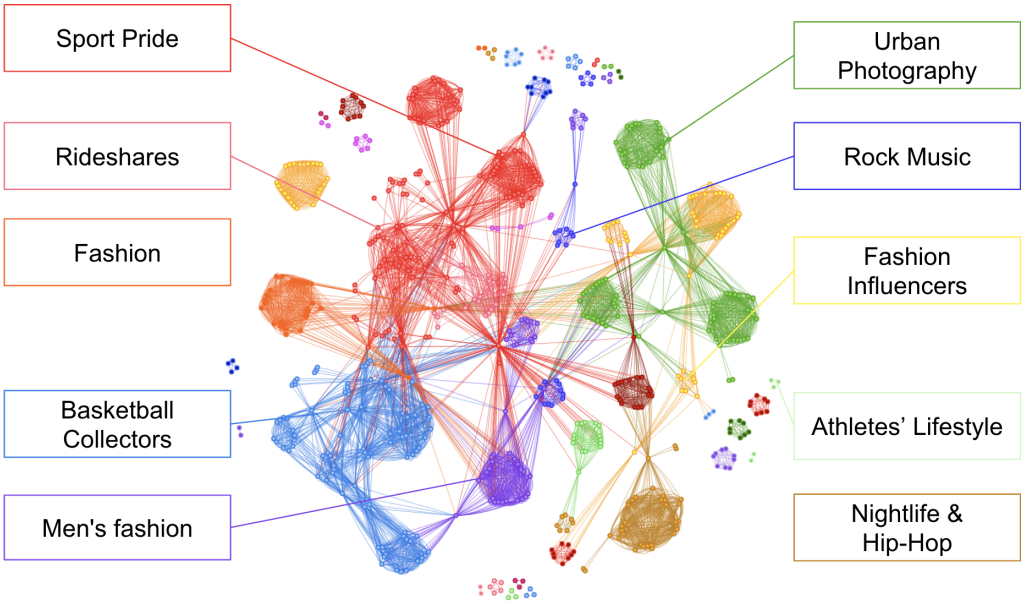

Figure 2.3: Community Detection Network of Nike in Cleveland.

2.4.1 Within-city and Between-brands

We begin with Nike and Adidas in San Francisco. Figure 2.4 presents the top ten communities for each brand, ranked by relevance. Relevance is determined by the proportion of images featuring the brand's logo that belong to a given community. Formally, the relevance of community C_i for brand b is measured as the ratio of images containing the logo of brand b that belong to the community C_i, over the total number of images featuring the logo of brand b. We also categorize each community into a broader macro-category.

[3]Selecting the top communities primarily addresses interpretability concerns. Nevertheless, we note that excluding residual clusters does not alter the overall insights derived from MBCA.

For interpretability, we group communities into broader macro-categories. For example, within the macro-category *Fashion*, we identify a *Streetwear & Sneakers* community (consumers embedded in sneaker culture and street style) and a *Fashion Influencers* community (consumers linking the brand to makeup, wardrobes, and outfit-of-the-day content).

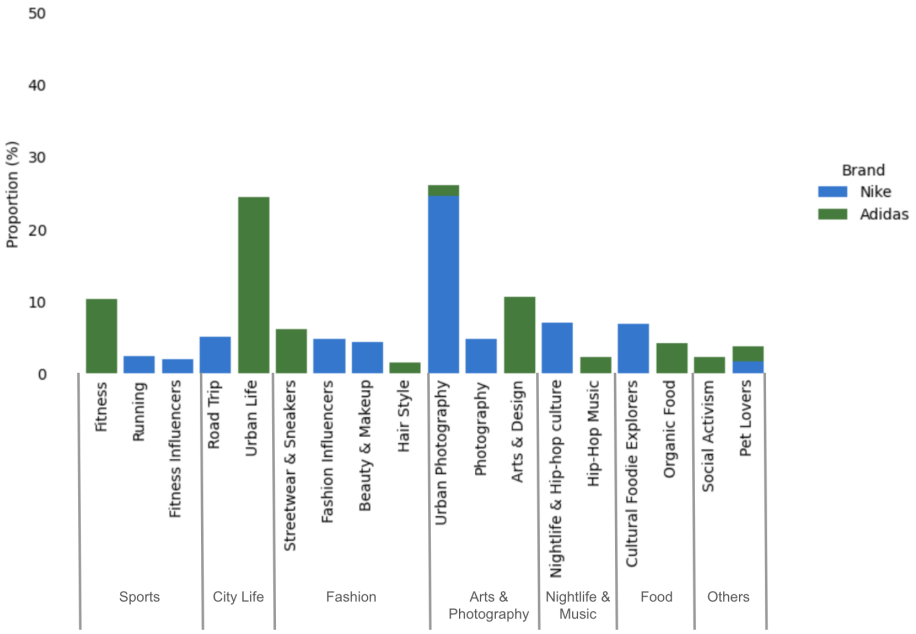

Notes. The Figure reports the proportion of each community in San Francisco.

Our analysis reveals both shared consumer interests and brand-specific distinctions. Nike and Adidas are both tied to sports-oriented consumers, but in different ways: Adidas communities emphasize running and fitness, while Nike connects more strongly with fitness influencers. In fashion, Adidas is rooted in streetwear and sneaker culture, whereas Nike leans toward influencers and beauty-related groups. Adidas is also prominent in *Arts & Design* communities, while Nike shows a stronger link to Photography. Both brands intersect with music and nightlife, but Adidas is more strongly associated with the hip-hop scene, while Nike associates more strongly with the club culture. Finally, unique brand-specific lifestyles emerge: Nike is connected to cultural foodie explorers, whereas Adidas resonates with organic food and social activism.

Taken together, the comparison shows that while both Nike and Adidas attract sports-oriented consumers, they also connect to a wide range of other lifestyle and cultural interests. MBCA captures these fine-grained, brand-specific differences, such as Adidas's ties to streetwear, art, and activism versus Nike's links to influencers, photography, and nightlife, providing a nuanced and multi-dimensional view of San Francisco's customer base. By mapping both shared and distinct communities, MBCA helps managers see how consumers connect brands to different habits and lifestyles. These insights support more precise targeting and allow strategies to reflect the cultural context of specific consumer groups.

2.4.2 Within-brand and Between-cities

We next analyze Adidas across San Francisco and Los Angeles. These two cities provide a strong test for our methodology: they are geographically close, belong to the same state, and share similar cultures. One might therefore expect highly similar brand communities. However, our results, shown in Figure 2.5, reveal meaningful differences, highlighting MBCA's ability to detect subtle variations in brand use across seemingly similar markets.

Several Adidas communities appear in both San Francisco and Los Angeles, such as *Urban Life*, *Streetwear & Sneaker*, and *Nightlife*; however, their relevance and meaning differ across the two cities. In San Francisco, Adidas is deeply tied to cultural and creative niches. *Urban Life* reflects nightlife and travel moments, while *Arts & Design* ties the brand in the city's artistic identity through contemporary art, interior design, and music subcultures. Other prominent groups include *Fitness*, *Streetwear & Sneakers*, and lifestyle communities such as *Organic Food*, *Social Activism*, and *Pet Lovers*. Creative communities like *Urban Photography* and *Hair Style* further illustrate how the brand integrates into San Francisco's everyday culture.

Los Angeles, by contrast, presents a more performance- and entertainment-oriented

communities. Here, *Nightlife* fuses fashion and club culture, while communities like *Yoga* echo the city's strong wellness ethos and *Cross-genre Music* spans DJs, festivals, and even global circuits. A dedicated *Sports Fan* community revolves around basketball and soccer, while fashion-related communities branch into *Streetwear & Sneakers*, drag-inspired *Beauty & Makeup*, and influencer-driven clusters such as *Fashion Influencers*. Notably, Los Angeles also connects Adidas to creative spaces such as *Pop-Comic Art* and *Auditions & Creative Hustle*, underscoring the city's relevant role as an entertainment and cultural hub.

Taken together, San Francisco positions Adidas through creativity, activism, and lifestyle experimentation, whereas Los Angeles associates the brand more strongly with performance, entertainment, and global style tribes. A parallel analysis of Nike communities in the two cities is reported in Appendix A1.1.

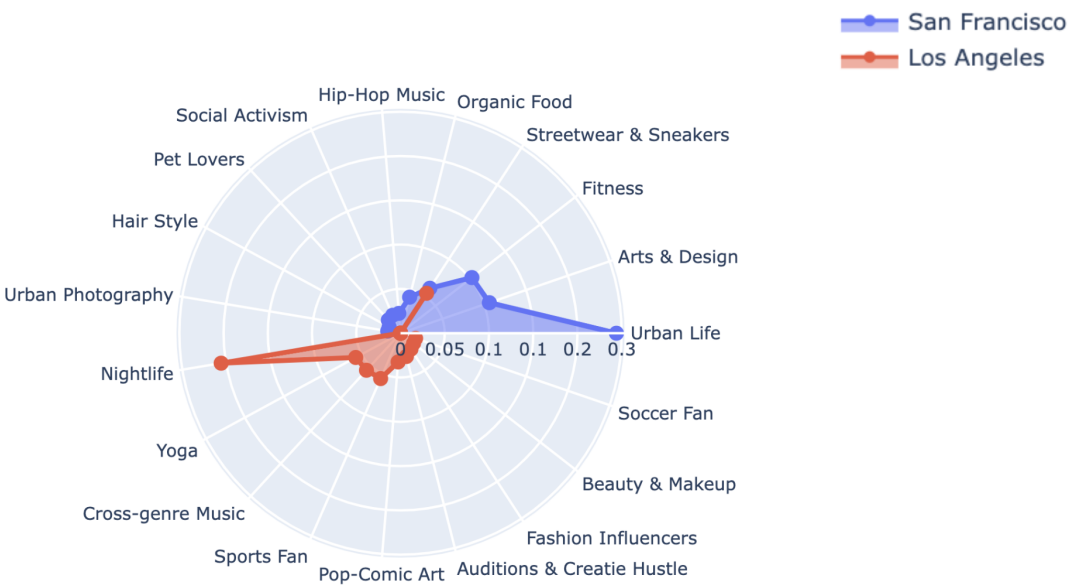

Figure 2.5: Adidas' Consumer Groups in San Francisco and Los Angeles.

After contrasting San Francisco and Los Angeles, we extend the analysis to six additional cities: Atlanta, Chicago, Cleveland, London, Miami, and Paris. Figure 2.6 shows the proportion of each community in each city using the same approach described before.

To facilitate comparisons, we organize communities into macro-categories (e.g., *Running* under *Sports*), which provides a consistent basis for evaluating consumer behavior across different locations. Reassuringly, this broader perspective reveals two consistent and systematic patterns: recurring themes that appear across multiple cities and distinct local expression tied to specific markets. For Adidas, *Arts*, *Fashion*, and *Sports* emerge consistently across all cities. We interpret these as the brand's core DNA, associations that marketers should emphasize in global campaigns (e.g., commercials or influencer collaborations) to sustain a cohesive worldwide identity. A parallel analysis for Nike communities is reported in Appendix A1.2. At the same time, MBCA highlights the value of local differentiation. City-specific communities surface in ways that carry direct marketing implications: Adidas connects with *Urban Life* in San Francisco, *Fashion* in London, and *Arts & Design* in Chicago. These insights are critical for managers looking for effective ways to balance global coherence with local relevance. By identifying both the enduring themes that form a brand's core DNA and the unique communities that reflect local culture, MBCA gives marketers a practical tool for balancing global consistency with local relevance.

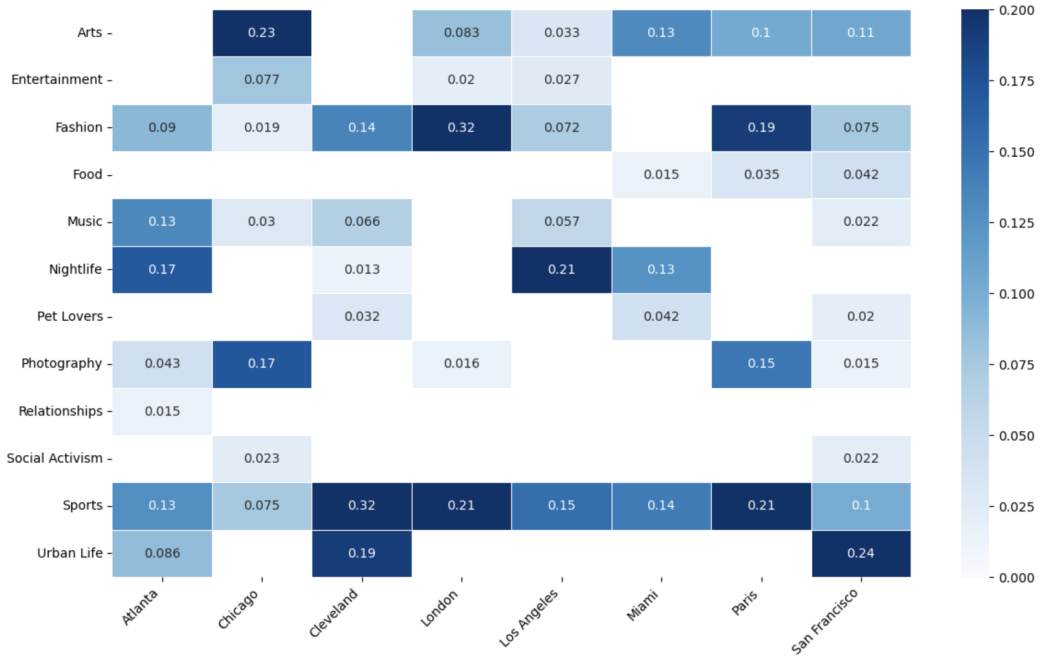

Note. The Figure reports the proportion of each community in the city. A darker color indicates a larger community.

Figure 2.6: Adidas' Communities Across Cities.

2.4.3 Brand Logo versus Brand Mention

As discussed the literature review section, analyses based solely on brand mentions risk providing an incomplete perspective on consumer brand use. Consumers often use a brand without explicitly tagging it, meaning that brand mentions capture only a subset of real interactions. To validate this conjecture, we compare the communities that emerge from images identified through visual logo detection with those identified through brand mentions. Figure 2.7 illustrates this comparison for Nike and Adidas in San Francisco.

The differences are striking. Communities derived from brand mentions are dominated by sports-related activities (21.9% of Adidas images and 43% for Nike images). By contrast, only a small share of images identified via visual logo detection fall into sports communities (10.2% for Adidas and 4.4% for Nike). This pattern reveals what we call *context bias*: consumers tend to mention a brand only when engaging in activities that

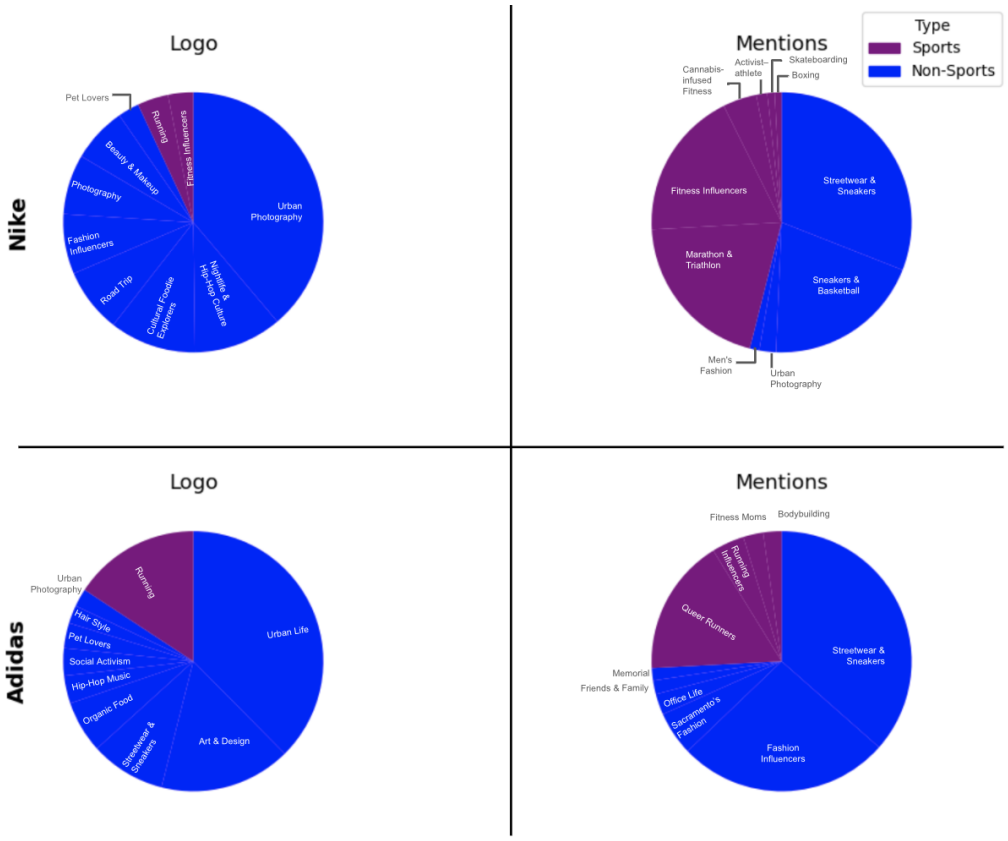

Figure 2.7: Logo versus Mention Comparison.

align with its core positioning (e.g., tagging Adidas in a fitness post). As a result, analyses restricted to brand mentions systematically understate the breadth of brand use.

Visual logo detection overcomes this limitation by capturing brand presence in a much wider range of everyday contexts. Consumers may wear Nike or Adidas while going out at night, attending art events, or engaging in casual social activities without ever mentioning the brand explicitly. These moments are invisible to text-based analysis but are central to understanding how brands actually live in consumers' daily lives. Prior research shows that brand interactions often occur subconsciously or in settings unrelated to a brand's official positioning (Ferraro et al., 2009). MBCA reveals implicit associations, giving marketers a more complete understanding of everyday brand use.

2.5 Conclusion

2.5.1 Theoretical Implications

We contribute to marketing theory in three main ways: demonstrating the value of multimodal approaches, offering a scalable method to capture local variation in brand use, and revealing structural biases in existing UGC research.

Advancing Multimodal Approaches to Consumer Insight

Most marketing research on UGC has relied on a single modality—numeric (Culotta and Cutler, 2016), textual (Lee and Bradlow, 2011; Netzer et al., 2012) or visual (Liu et al., 2020; Zhang and Luo, 2023). Our study contributes to the growing stream of research on multimodal marketing analysis (Klostermann et al., 2018; Grewal et al., 2021; Hartmann et al., 2021) by showing how visual and textual content can be integrated to generate deeper insights into brand meaning. Specifically, we use computer vision to detect brand logos, thereby capturing brand presence regardless of mention, and combine this with network-based analysis of co-occurring hashtags to infer and interpret the cultural and

situational context of brand use. This approach enables researchers to go beyond either modality and to avoid treating either in isolation.

This integration matters theoretically. Images capture brand use that is unprompted and often unspoken, while hashtags supply the semantic cues that situate this use. By bringing these signals together, we obtain both an accurate representation of consumer behavior and an interpretable mapping of how brands are embedded in culture. The contribution to multimodal research is thus twofold: we demonstrate that integration is technically feasible, and we show it is substantively meaningful. This combined lens enables new avenues for studying brand meaning, community formation, and cultural variation through the combined lens of what consumers show and how they frame what they show. We believe that our multimodal approach better captures and reflects the complexity of consumer expression in contemporary digital environments.

A New Scalable Framework for Studying Localized Brand Meaning

Research on global branding has traditionally examined standardized versus localized positioning across countries, focusing on how global brands navigate differences in language, cultural values, or institutions (Alden et al., 1999; Steenkamp, 2001). Similarly, most UGC-based studies aggregate content globally, producing a single representation of brand associations (Klostermann et al., 2018; Hartmann et al., 2021; Rust et al., 2021).

Both approaches tend to overlook within-country variation, particularly in large, culturally diverse markets. Our findings highlight the importance of this neglected level of analysis. While some themes recur across geographies, what we term a brand's DNA, we also observe systematic city-level differentiation. The same brand may be associated with nightlife in one city, fitness in another, and food culture in a third. These differences persist even across proximate cities such as San Francisco and Los Angeles. Thus, brand meaning is not monolithic but emerges from localized consumer practices that vary across urban contexts.

Scholars have documented such local variation through qualitative methods such as ethnography and in-depth interviews. For example, Thompson and Arsel (2004) describe in great detail how Starbucks is localized and contested in North American cities, but their insights rely on immersive fieldwork that is difficult to scale. Consequently, city-level branding remains a theoretically significant but methodologically underdeveloped area.

We address this gap by introducing a scalable, data-driven framework for analyzing how global brands acquire local meaning. Our methodology allows researchers to empirically investigate the urban embedding of brands across multiple markets without relying exclusively on resource-intensive qualitative work. In doing so, we provide a new tool for studying *glocalization*: one that preserves sensitivity to cultural specificity while enabling systematic comparison across diverse urban markets.

Revealing Context Bias in UGC Research

Much of the UGC literature analyzes data that consumers actively provide. Studies typically examine brand mentions in product reviews or tweets (Tirunillai and Tellis, 2014; Liu et al., 2019; Timoshenko and Hauser, 2019), or hashtags on social platforms or images (Nam et al., 2017; Klostermann et al., 2018). These approaches yield valuable insight into how consumers articulate their preferences, but they inherently privilege intentional expressions (i.e., moments when consumers deliberately position the brand at the center of their message). In effect, they capture what consumers choose to say about a brand.

Our findings document a structural limitation of this approach, which we term context bias: the systematic overrepresentation of brand use in brand-aligned contexts and the underrepresentation of everyday, incidental uses. Put differently, consumers disproportionately tag a brand when it aligns with its canonical purpose (e.g., mentioning Nike in a workout post), while ignoring more mundane situations where the brand is present but not focal (e.g., wearing Nike at a casual dinner).

By documenting this context bias, we extend UGC research from declarative expressions—what consumers choose to say (Toubia and Netzer, 2017)—to incidental brand presence—how consumers live with brands and how they curate their online identities. This reframes UGC as a continuum of intentional and unintentional interactions. Whereas prior research has largely examined explicit, deliberate brand mentions, our findings emphasize the background role brands play in everyday life. This perspective invites scholars to reconceptualize UGC not just as a record of consumer speech but also as a set of behavioral traces that capture the unspoken (i.e., invisible) embedding of brands in daily routines.

2.5.2 Managerial Implications

Our study offers three key managerial implications: marketers can use MBCA to translate consumer insights from UGC into scalable, context-aware execution, localize campaigns at the city level, and uncover unprompted brand use.

Translating UGC-Based Insight Into Execution

van den Driest et al. (2016) contend that developing customer insights is not sufficient; rather, the real competitive advantage for brands lies in their ability to translate those insights into creative and strategic execution. As the authors put it, leading firms must become "insights engines" that move beyond passive reporting to proactively develop marketing campaigns and drive execution across functions.

MBCA outputs consumer communities defined by semantically coherent hashtags, which managers can use as structured prompts for large language and image generation models to efficiently generate campaign concepts, taglines, and visual assets that reflect the specific language and interests of each community. In this way, MBCA creates a streamlined pipeline from UGC to creative production, reducing the risk of miscommunication between research and creative teams, and enabling more coherent, responsive, and

culturally aligned campaign development.

To illustrate how managers can apply this workflow in practice, we use GPT-5 to generate marketing campaign concepts for Nike shoe launch in San Francisco, drawing on three communities identified by MBCA: *Urban Photography*, *Cultural Foodie Explorers*, and *Beauty & Makeup*. By inputting the community labels and hashtags, we obtain fully formed campaign ideas. We report the full prompts in Appendix A2 and present the resulting campaigns in Appendix A3. This example illustrates how manages can translate the output of MBCA into fully formed campaign strategies with minimal translation effort between research and creative teams. It highlights the strategic value of our method not only in retrieving consumer insight but in enabling faster, more scalable, and locally relevant marketing execution, a core challenge for global brands seeking both resonance and reach.

Concluding, Table 2.2 contrasts MBCA with other market research techniques. While existing methods excel in particular domains, MBCA stands out by combining scale with contextual richness and by addressing context bias. This unique positioning underscores MBCA's unique value as a bridge between quantitative tracking, cultural interpretation, and real-time digital monitoring.

DimensionMethod	Qualitative	Quantitative	Traditional UGC Analyses	MBCA
Data Source	Fieldwork, interviews, focus groups	Surveys	Hashtags, captions, tagged images	Logo detection + hashtags
Expression Captured	Everyday practices	Stated attitudes	Explicit brand talk	Explicit + incidental use
Context Coverage	Rich local context	Limited to questions	Brand-aligned contexts	Full range, including everyday
Scale	Narrow (dozens)	Broad (hundreds–thousands)	Broad (millions)	Very broad (hundreds of millions)
Geographic Resolution	Local, micro-level	National or regional	Typically global/national	City-level, cross-city mapping
Bias Exposure	Researcher interpretation	Self-report, desirability	*Context bias* (mentions only)	Reduces context bias (logos + incidental use)
Managerial Insight	Cultural depth, identity meaning	Equity benchmarks, awareness	Reputation monitoring, association maps	Consumer communities, unprompted uses, global DNA vs. local variation, campaign inputs

Table 2.2: Comparing MBCA to Market Research Techniques.

Balancing Insight Depth and Geographic Reach

Marketers have long recognized the importance of adapting campaigns and products across countries to account for cultural differences. Global brands such as Coca-Cola, Nike, and Heineken routinely tailor messaging and offerings to local consumer expectations. For instance, in Brazil, Heineken launched the "Uncap the Brazilian Soul" campaign, celebrating the country's vibrant culture during the Rio Carnival, while Nike varies its campaign tone across continents, emphasizing athleticism in the United States, fashion and street culture in Europe, and collective achievement in East Asia. These efforts reflect a well-established practice: adjusting global strategy to national cultural contexts.

However today, national-level adaptation is often insufficient. Luxury brands, for instance, now treat city-level marketing as a strategic imperative designing pop-ups, collaborations, and campaigns that speak directly to the cultural DNA of individual cities (Chitrakorn, 2022). These initiatives reflect a growing recognition that meaningful brand engagement increasingly requires more than national-level adaptation, especially in large and culturally heterogeneous markets, such as the United States, China, or Brazil, where consumer lifestyles and brand behaviors often diverge sharply across cities.

However, designing resonant campaigns at the city level presents practical challenges. Nike's "Nothing Beats a Londoner" campaign illustrates both the impact and the cost of hyper-local strategy. The campaign stemmed from a widely shared concern that "Nike had lost touch with the real kids of London and our campaigns didn't talk to them in their own language anymore".[4] To reconnect with urban youth, the creative team reportedly "talked to hundreds of kids to find out who their London heroes are in sport, in music, in real life." The effort produced a widely celebrated campaign, with Nike searches in London increasing by 93% in its aftermath. However, replicating this level of ethnographic depth across dozens of cities would be prohibitively time- and resource-intensive, limiting its scalability as an ongoing strategic model. By contrast, Heineken's "Cities of the World"

[4]For further detail see: https://shorturl.at/cKpsj.

campaign represented a more scalable, but less culturally embedded, approach to city-level engagement.[5]

Targeting six global cities (New York, Shanghai, Berlin, Amsterdam, London, and Rio de Janeiro), the brand released limited-edition bottles named after each location and invited consumers to "unlock the secrets of their city" through a uniform, globally consistent creative platform. Rather than embedding itself deeply in any one city context, the campaign leveraged universal themes of exploration, mobility, and cosmopolitan identity, thereby enabling replication across markets without the need for extensive local research.

These two campaigns clearly highlight the strategic trade-off between local depth and executional scalability. MBCA helps reconcile this trade-off by offering a scalable, data-driven approach to city localization. Our method uncovers city-specific consumer patterns that can guide creative strategy without the cost of bespoke ethnography. This enables marketers to develop campaigns that are both locally resonant and operationally feasible, advancing the strategic promise of city-level marketing beyond one-off successes toward a repeatable and scalable capability.

Seeing What Consumers Do Not Say

Traditional tools for assessing brand perception, such as focus groups, surveys, and even UGC tracking, focus primarily on declared interactions, such reviews, ratings, or brand mentions. While useful, these methods systematically miss how brands appear in the background of everyday life, when consumers incorporate them into routines without calling attention to them.

MBCA addresses this limitation by providing a scalable way to track how brands are actually used across diverse contexts. For managers, this delivers a more holistic view of brand use that reflects both consumers' expressed attitudes and their actual behaviors. Strategically, this capability unlocks several advantages. For instance, it enables early

[5]For further detail see: https://shorturl.at/YMQMN.

detection of emerging usage patterns before they surface in traditional data, informs product innovation and campaign themes, facilitates influencer partnerships, and allows for continuous tracking of brand integration into consumer lifestyles. Most importantly, it shifts insight from "episodic snapshots" to "ongoing monitoring", moving brands from reactive adaptation toward proactive, context-aware strategy.

2.5.3 Limitations

Our study has limitations that open promising directions for future research. First, future research is needed to examine how the magnitude and nature of context bias may vary across categories and cultural settings. For instance, brands with narrow usage domains, such as sports or luxury, may exhibit stronger bias than brands with broader lifestyle associations, such as technology or food. Similarly, cross-cultural comparison could reveal whether local norms intensify or mitigate the selective tagging of brands. By developing context bias as a construct, we invite subsequent studies to explore these boundary conditions and to refine the understanding of when and how UGC provides a representative view of consumer–brand engagement.

Second, our analysis draws exclusively on Instagram, the leading image-sharing platform during the study period. While Instagram is well-suited for multimodal analysis given its emphasis on visual and textual content, consumer expression varies across platforms (e.g., TikTok's short-form video). Future work could extend MBCA to other platforms to assess whether context bias and city-level variation in brand meaning generalize across modalities of consumer expression. Such extensions would enrich understanding of how platform features influence consumer–brand interactions.

Third, our dataset covers 2015–2018, a period that allows us to observe brand use across a stable pre-pandemic context. However, consumer–brand interaction is dynamic, and future research could examine longitudinal changes in brand communities. A temporal extension would allow scholars to examine not only cross-sectional differences but

also how brand DNA and local expressions evolve over time. Fourth, we focus on two athletic brands (Nike and Adidas) to illustrate the method's application. This focus ensures comparability across firms while spanning different positions in the brand life cycle. Nonetheless, consumer expression may differ for luxury, fast-moving consumer goods, or technology brands, where cultural meanings and contexts of use are distinct. Applying MBCA to other categories would test the generalizability of context bias and may reveal new forms of brand–consumer interaction specific to different industries. Fourth, our analysis centers on eight major U.S. and European cities. This selection balances computational feasibility with diversity in cultural contexts. Future work could apply MBCA in emerging markets or smaller urban contexts to examine whether local brand communities exhibit similar structures. Such studies would deepen understanding of glocalization at multiple scales, from global metropolises to regional hubs.

A1 Appendix: Between-city and Within-brand

A1.1 Nike Communities in San Francisco and Los Angeles

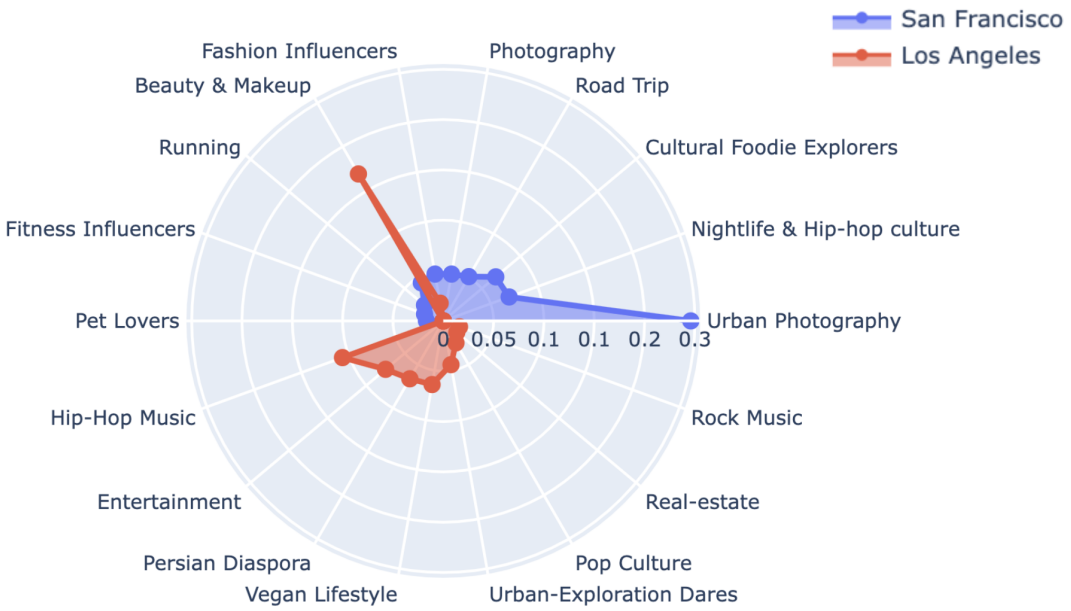

Figure 8: Nike's Consumer Groups in San Francisco and Los Angeles.

Figure 8 reports the results for Nike, comparing San Francisco and Los Angeles. In San Francisco, Nike appears in creative and lifestyle clusters such as *Urban Photography,* where street art, food culture, and design tags mix with city landmarks, and *Nightlife & Hip-hop culture,* anchored in local music promotion and club energy. Consumers also express Nike through *Cultural Foodie Explorers, Road Trip,* and photography-heavy communities, alongside fashion-driven clusters like *Fashion Influencers* and *Beauty & Makeup.* Fitness-oriented groups are present as well, ranging from *Running* logs to *Fitness Influencers,* complemented by lifestyle elements such as *Pet Lovers.* By contrast, Los Angeles communities place Nike in entertainment and performance contexts. Alongside *Beauty & Makeup* and *Hip-hop Music,* we observe a strong *Entertainment* cluster tied to Hollywood and the

music industry, as well as lifestyle niches such as *Vegan Lifestyle* and *Urban-Exploration Dares.* LA also features influencer and fashion-driven groups (*Fashion Influencers, Entertainment,* and *Real-estate*) and music subcultures such as *Rock Music.* Taken together, Nike's San Francisco presence highlights creativity, food, and urban culture, while its Los Angeles communities connect the brand more strongly to entertainment, celebrity, and lifestyle expression.

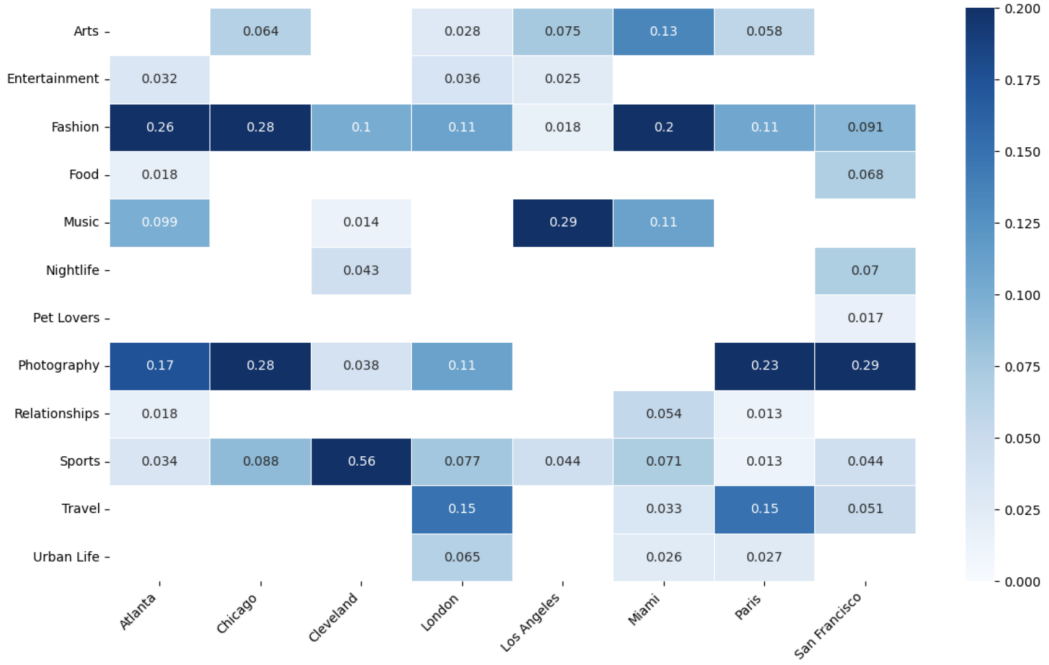

Note. The Figure reports the proportion of each community in the city. A darker color indicates a larger community.

Figure 9: Nike's Communities Across Cities.

A1.2 Nike Communities Across Cities

Extending the analysis to six additional cities—Atlanta, Chicago, Cleveland, London, Miami, and Paris—confirms the dual pattern of recurring themes and local variation. Figure 9 shows the proportion of Nike communities in each city. Several themes recur across locations, including *Fashion*, *Photography*, and *Sports*, which together form part of Nike's brand DNA. At the same time, distinctive local communities emerge: in San Francisco, Nike is strongly tied to photography; in Los Angeles, it connects to music and celebrity-driven lifestyles; in Atlanta, photography and community engagement are more salient; and in European cities such as London and Paris, fashion and global streetwear dominate. These results highlight that Nike's global identity is consistently rooted in fashion, photography, and sports, while its local expressions vary significantly across urban contexts.

A2 Appendix: Prompt Documentation

This appendix contains the sequence of prompts used to generate the campaign concepts and ads. The interaction was structured as a progressive briefing to simulate a creative collaboration with a marketing strategist. Each prompt built upon the previous one to elicit increasingly specific and detailed campaign responses.

Prompt 1: Role Definition *You are a marketing expert. You create campaigns for global brands.*

Prompt 2: Scenario Context *You work for Nike, which wants to engage with its consumers in San Francisco. You have conducted some preliminary analysis to better understand who they are. First, you have collected all Instagram images that contain the Nike logo posted by consumers living in San Francisco. Then, you have analyzed the hashtags contained in these pictures to identify communities of consumers based on their preferences, interests, and lifestyles.*

Prompt 3: Campaign Direction *I like them. Now, develop an ad for each community to support the different campaigns. The hashtag should contain a slogan and an image of the new shoes. You can adapt the background, color, and content of the image to get the best fit with the specific campaign targeting the community.*

Prompt 4: Visual Request *Generate visual mockups of these ads.*

A3 Appendix: Community Campaigns

Component	Urban Photography	Foodie Explorers	Beauty & Makeup
Key Insight	Creators capture San Francisco as a canvas — blending landmarks, street art, coffee culture, sports pride, and design into a visual lifestyle.	Culture-hunters are always chasing new flavors, neighborhoods, and nightlife.	Beauty is how you feel: wellness, self-expression, and clean choices matter.
Shoe Positioning	Performance-driven yet photogenic sneaker, built for movement across the city and styled to stand out in every shot.	Lightweight, cushioned sneaker built for day-to-night city life and endless walking.	Sleek minimalist design, skin-tone inspired palettes, sustainable materials, and all-day comfort.
Campaign Tactics	• **Instagram Challenge: "#ThroughTheCitySF"** — Community invited to post Nike-inspired city shots (murals, cable cars, sports, everyday walks). Featured creators showcased on Nike SF. • **Local Collab Drops** — Limited-edition sneakers inspired by SF street art, designed with Mission District muralists and shot by local IG photographers. • **Coffee x Creativity Pop-Ups** — Nike x Blue Bottle/Tartine cafés with sneaker showcases, live mural sessions, and design talks (with firms like frog/IDEO). • **Sports Fandom Hook** — Warriors/Giants game-day creator activations; Nike visuals featured in stadium, winners get behind-the-scenes passes.	• **"Bites & Kicks" Map** — Co-created food/fashion walking map with local influencers. Five food spots + five photo stops. QR/Nike app access with check-in rewards. • **Night Market Pop-Up** — Chinatown or North Beach event with curated local food stalls; free for those wearing the shoe. • **UGC Contest: #TastesLikeSF** — Users pair their shoes with local dishes; winners receive fictional Nike x Eatery sneaker collab. • **In-App Content Series** — "24 Hours in SF" influencer-led food tour, shot while wearing the sneaker.	• **Influencer Series: GRWM "Feet First"** — Beauty creators incorporate the sneaker into Get Ready With Me routines that show glam meets lifestyle. • **Nike x Clean Beauty Pop-Up** — Downtown activation (e.g., Westfield/ Fillmore) with sneaker try-ons, natural makeup booths, demos by local brands (e.g., Credo Beauty). • **#StepIntoBeauty Campaign** — Users share clean beauty + sneaker looks. Winners get Nike wellness kits (skincare, sneakers). • **Colorway Naming Activation** — Community votes on sneaker shade names inspired by makeup tones (e.g., "Rosé Clay", "Matte Mocha").

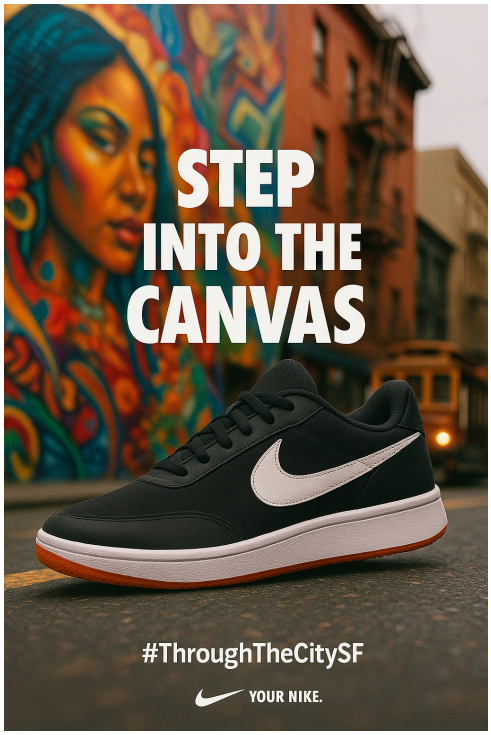

((a)) Urban Photography.

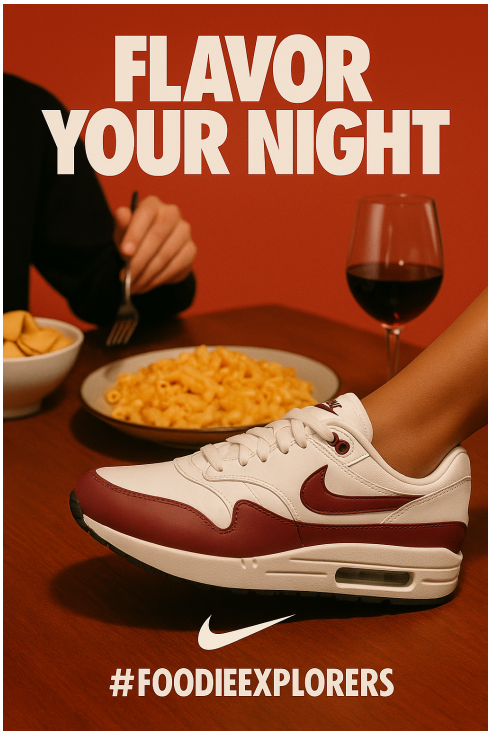

((b)) Foodie Explorers.



((c)) Beauty & Makeup.

Figure 10: GPT-5 Generated ads to target each community.

References

Alden, D. L., Steenkamp, J.-B. E., and Batra, R. (1999). Brand positioning through advertising in asia, north america, and europe: The role of global consumer culture. *Journal of Marketing*, 63(1):75–87.

Chitrakorn, K. (2022). Global luxury brands go deeply local. suburbia isn't far off.

Chung, J., Johar, G. V., Li, Y., Netzer, O., and Pearson, M. (2022). Mining consumer minds: Downstream consequences of host motivations for home-sharing platforms. *Journal of Consumer Research*, 48(5):817–838.

Culotta, A. and Cutler, J. (2016). Mining brand perceptions from Twitter social networks. *Marketing Science*, 35(3):343–362.

Dang, C., Kwan, C. M., Jia, J. S., and Shi, Y. (2025). Express: When words meet visuals: How content composition drives social media engagement for marketer-generated content. *Journal of Marketing Research*, page 00222437251373042.

Deng, J., Dong, W., Socher, R., Li, L.-J., Li, K., and Fei-Fei, L. (2009). Imagenet: A large-scale hierarchical image database. In *2009 IEEE Conference on Computer Vision and Pattern Recognition*, pages 248–255. Ieee.

Dutta, A. and Zisserman, A. (2019). The via annotation software for images, audio and video. In *Proceedings of the 27th ACM International Conference on Multimedia*, pages 2276–2279.

Ferraro, R., Bettman, J. R., and Chartrand, T. L. (2009). The power of strangers: The effect of incidental consumer brand encounters on brand choice. *Journal of Consumer Research*, 35(5):729–741.

Fournier, S., Quelch, J. A., and Rietveld, J. (2016). To get more out of social media, think like an anthropologist. *Harvard Business Review*.

Golino, H. F. and Epskamp, S. (2017). Exploratory graph analysis: A new approach for estimating the number of dimensions in psychological research. *PloS one*, 12(6):e0174035.

Grewal, R., Gupta, S., and Hamilton, R. (2021). Marketing insights from multimedia data: Text, image, audio, and video. *Journal of Marketing Research*, 58(6):1025–1033.

Hartmann, J., Heitmann, M., Schamp, C., and Netzer, O. (2021). The power of brand selfies. *Journal of Marketing Research*, 58(6):1159–1177.

Hollow, J. (2023). Social listening has a long way to go. https://researchworld.com/articles/social-listening-has-a-long-way-to-go.

Jagabathula, S., Subramanian, L., and Venkataraman, A. (2018). A model-based embedding technique for segmenting customers. *Operations Research*, 66(5):1247–1267.

Klostermann, J., Plumeyer, A., Böger, D., and Decker, R. (2018). Extracting brand information from social networks: Integrating image, text, and social tagging data. *International Journal of Research in Marketing*, 35(4):538–556.

Kushner, D. (2025). The overlooked edge: Content intelligence meets social listening. *Forbes*.

Lee, T. Y. and Bradlow, E. T. (2011). Automated marketing research using online customer reviews. *Journal of Marketing Research*, 48(5):881–894.

Liu, L., Dzyabura, D., and Mizik, N. (2020). Visual listening in: Extracting brand image portrayed on social media. *Marketing Science*, 39(4):669–686.

Liu, X., Lee, D., and Srinivasan, K. (2019). Large-scale cross-category analysis of consumer review content on sales conversion leveraging deep learning. *Journal of Marketing Research*, 56(6):918–943.

Moe, W. W., Netzer, O., and Schweidel, D. A. (2017). Social media analytics. *Handbook of Marketing Decision Models*, pages 483–504.

Nam, H., Joshi, Y. V., and Kannan, P. K. (2017). Harvesting brand information from social tags. *Journal of Marketing*, 81(4):88–108.

Nam, H. and Kannan, P. K. (2014). The informational value of social tagging networks. *Journal of Marketing*, 78(4):21–40.

Netzer, O., Feldman, R., Goldenberg, J., and Fresko, M. (2012). Mine your own business: Market-structure surveillance through text mining. *Marketing Science*, 31(3):521–543.

Pons, P. and Latapy, M. (2005). Computing communities in large networks using random walks. In *Computer and Information Sciences-ISCIS 2005: 20th International Symposium, Istanbul, Turkey, October 26-28, 2005. Proceedings 20*, pages 284–293. Springer.

Redmon, J. and Farhadi, A. (2018). Yolov3: An incremental improvement. *arXiv preprint arXiv:1804.02767*.

Rust, R. T., Rand, W., Huang, M.-H., Stephen, A. T., Brooks, G., and Chabuk, T. (2021). Real-time brand reputation tracking using social media. *Journal of Marketing*, 85(4):21–43.

Schwarz, R. (2023). 10 commandments of user-generated content. https://www.forbes.com/councils/forbescommunicationscouncil/2023/03/21/10-commandments-of-user-generated-content/.

Steenkamp, J.-B. E. (2001). The role of national culture in international marketing research. *International Marketing Review*, 18(1):30–44.

Thompson, C. J. and Arsel, Z. (2004). The Starbucks brandscape and consumers'(anticorporate) experiences of glocalization. *Journal of Consumer Research*, 31(3):631–642.

Timoshenko, A. and Hauser, J. R. (2019). Identifying customer needs from user-generated content. *Marketing Science*, 38(1):1–20.

Tirunillai, S. and Tellis, G. J. (2014). Mining marketing meaning from online chatter: Strategic brand analysis of big data using latent dirichlet allocation. *Journal of Marketing Research*, 51(4):463–479.

Toubia, O. and Netzer, O. (2017). Idea generation, creativity, and prototypicality. *Marketing science*, 36(1):1–20.

Ultralytics (2021). Yolov5: A state-of-the-art real-time object detection system. https://docs.ultralytics.com.

van den Driest, F., Sthanunathan, S., and Weed, K. (2016). Building an insights engine. *Harvard business review*, 94(9):64–74.

Wang, X., He, J., Curry, D. J., and Ryoo, J. H. (2022). Attribute embedding: Learning hierarchical representations of product attributes from consumer reviews. *Journal of Marketing*, 86(6):155–175.

Yang, Y., Zhang, K., and Kannan, P. (2022). Identifying market structure: A deep network representation learning of social engagement. *Journal of Marketing*, 86(4):37–56.

Zhang, M. and Luo, L. (2023). Can consumer-posted photos serve as a leading indicator of restaurant survival? Evidence from Yelp. *Management Science*, 69(1):25–50.

Chapter 3

Assessing Entrepreneurial Activity through Deep Learning: A Geographic Analysis Using Instagram Data

Francesco Grossetti[1], **Paola Cillo**[2], **Margherita Caprara**[3], **Gaia Rubera**[3]

[1]Department of Accounting and BIDSA, Bocconi University

[2]Department of Management and Technology, Bocconi University

[3]Department of Marketing, Bocconi University

This study investigates the relationship between narcissistic self-presentation on social media and entrepreneurial activity across urban areas. Drawing on a novel dataset of over 5 million Instagram images from 33 major cities between 2015 and 2018, we use a deep learning model (YOLO v3) to detect selfies and employ them as behavioral proxies for non-pathological narcissism. We match this data with longitudinal records of startup activity from Crunchbase and entrepreneurial quality data from the Startup Cartography Project. Our preliminary results reveal a positive and significant association between selfie prevalence and the quantity of startups founded at the city level. However, we find no evidence that selfie prevalence correlates with higher startup quality, as measured by the Entrepreneurial Quality Index (EQI). These findings suggest that narcissistic self-presentation may spur entrepreneurial activity, but not necessarily its success or quality, highlighting the complex role of personality traits in shaping regional economic outcomes.

Keywords: Entrepreneurship; user-generated content; computer vision; artificial intelligence

3.1 Introduction

Narcissistic self-presentation is increasingly visible in the digital sphere, where platforms like Instagram enable individuals to curate and share images of themselves with wide audiences. Selfies, which are self-portraits typically shared on social media, have emerged as a common and accessible mode of this visual self-expression. Scholars have identified narcissism as a key predictor of selfie posting frequency, photo editing behaviors, and broader self-presentation strategies online (Fox and Rooney, 2015; Barry et al., 2017). These behaviors are not merely superficial; rather, they represent deeper psychological orientations toward self-importance, attention-seeking, and status signaling. Narcissism has also been associated with professional ambition, leadership development, and, important for this study, entrepreneurial tendencies (Leung et al., 2021; Engelen et al., 2016).

Entrepreneurship is often portrayed as a high-risk, high-reward domain that demands confidence, visibility, and a tolerance for ambiguity, traits that overlap considerably with the narcissistic personality profile (Navis and Ozbek, 2017). Recent experimental evidence by Yadav and Batra (2023) supports this connection, showing that narcissism significantly increases entrepreneurial intentions by enhancing self-efficacy, perceived behavioral control, and risk propensity. However, despite growing scholarly interest, most studies focus on individual-level intentions or CEO-level traits, and few investigate whether narcissistic behavior translates into observable economic activity. Moreover, virtually no research has examined how narcissistic tendencies expressed through public social media behaviors, such as selfies, relate to entrepreneurial outcomes at the city level. Prior research has shown that individual personality traits cluster geographically and can predict economic outcomes such as entrepreneurship (Obschonka, 2017). These geographic psychological patterns are robust across estimation strategies and spatial controls (**?**), supporting the view that aggregated personality traits can meaningfully influence regional behavior. This paper addresses this gap by combining computer vision analysis of over 5 million Insta-

gram images across 33 U.S. cities with longitudinal data on startup activity and startup quality. We find that narcissistic self-presentation, proxied by selfie prevalence, is positively associated with the quantity of startups but unrelated to their quality, as measured by the Entrepreneurial Quality Index (EQI) (Andrews et al., 2022). This suggests that while narcissism may catalyze entrepreneurial action, it does not necessarily enhance entrepreneurial performance or impact.

3.2 Theoretical Background

Narcissism has long been a focus in psychology, first introduced by the early work of Ellis (1898) who borrowed the label from the myth of Narcissus - a young man who became so enamored with his own reflection that he ultimately drowned. Narcissism is defined as a stable personality trait (Engelen et al., 2016) characterized by an inflated sense of self-importance and superiority over others. Narcissism exists on a continuum within the general population and can manifest in behaviors that are relevant to leadership and entrepreneurial activity. Research has shown that narcissistic traits, such as confidence, risk-taking, and a desire for admiration, can drive individuals to seek out positions of influence, such as leadership roles or entrepreneurial ventures (Chatterjee and Hambrick, 2007). This desire for public recognition and achievement aligns closely with entrepreneurial orientation, where narcissistic individuals are more likely to engage in high-risk, high-reward behaviors (Campbell et al., 2004), such as starting new ventures.Therefore, narcissism is not just a trait linked to personal self-perception but also has broader implications for economic activity, particularly in entrepreneurial settings.

The rise of big data has opened new avenues for studying how personality traits like narcissism influence regional entrepreneurial behavior. Publicly available digital data, such as social media posts, provides valuable insights into the psychological dimensions of individuals and groups (Obschonka et al., 2020). As social media enables people to present

themselves more readily and easily than ever (Belk, 2013), users' choices in sharing content and engaging with others offer a unique opportunity to observe behaviors, attitudes, and personality traits that, ultimately, can be linked to entrepreneurial outcomes. For instance, research has shown that traits like extraversion, openness, and conscientiousness are linked to entrepreneurial culture in a specific region, with areas displaying higher concentrations of these traits exhibiting increased startup activity (Obschonka, 2017).

Traditionally, research on personality traits, including narcissism, has depended on self-reported data collected through personality questionnaires (Brandstätter, 2011; Obschonka, 2017). However, this method is both time-consuming and prone to selection biases, particularly when attempting to capture data on a local cultural level. These challenges have driven the search for more scalable and objective methods for measuring personality traits (Gebru et al., 2017). Computational approaches — particularly those based on machine learning and deep learning — offer scalable and objective tools to infer psychological traits from digital traces. For instance, Adjerid and Kelley (2018) highlights how machine learning can complement traditional psychological methods, enabling the extraction of behavioral signals from large-scale digital data.

In the case of narcissism, digital behaviors like posting selfies on social media offer a proxy for measuring narcissistic tendencies (Barry et al., 2017; Weiser, 2015). Narcissists, known for their inflated self-views and constant need for admiration (Campbell and Foster, 2011), frequently use platforms like Instagram to present themselves and seek validation. Selfie-posting behavior has been linked to narcissism, with individuals who exhibit higher narcissistic tendencies more likely to post more and more curated selfies as a form of self-promotion (Fox and Rooney, 2015; Sorokowski et al., 2015). By leveraging selfies as an indicator of narcissistic traits, this study uses publicly available big data to assess how narcissism at the regional level correlates with entrepreneurial activity, specifically the number of startups founded.

While narcissistic traits are associated with greater entrepreneurial orientation and an

increased likelihood of starting ventures, their impact on the quality of these startups is more nuanced. Research on firm performance suggests that narcissistic leaders may drive innovation and take bold risks that benefit firms' performance (Chatterjee and Hambrick, 2007). However, their self-centeredness and difficulty in managing teams effectively can lead to poor long-term performance (de Vries and Miller, 1985). Similarly, Engelen et al. (2016) found that while narcissistic CEOs may thrive in dynamic markets, their leadership often weakens the link between entrepreneurial orientation and performance in more stable environments. Thus, while narcissistic traits can drive the formation of startups, they may also challenge the long-term success and sustainability of these ventures.

3.3 Data and Method

This study integrates three main data sources. First, we developed a web scraping tool to collect Instagram images and associated metadata from 2015 to 2018 across 33 major cities worldwide. Cities include urban hubs in the U.S. (e.g., New York, San Francisco), Europe (e.g., Berlin, Milan, London), Asia (e.g., Shanghai), South America (e.g., São Paulo), and Oceania (e.g., Sydney). Instagram, as of 2023, is the world's leading image-sharing platform, with approximately 2 billion monthly users and 100 million daily photo uploads (Dixon, 2023). Our scraper sampled images every 10 minutes from each city's geographic location, ultimately gathering over 2 terabytes of visual data.

Second, to assess entrepreneurial quantity, we use Crunchbase data, capturing the number of startups founded in each city per month.

Third, we employ data from the Startup Cartography Project (Andrews et al., 2022), which estimates the expected growth potential of startups across U.S. ZIP codes. We use their Entrepreneurial Quality Index (EQI), which reflects the average quality of new ventures in a given location and year, focusing on data from 2015 and 2016 due to overlap constraints with our Instagram and Crunchbase data.

3.3.1 Selfie Detection

To identify selfies in the dataset, we implemented a two-step classification pipeline based on object detection and facial framing. We used YOLOv3 (You Only Look Once v3), a state-of-the-art convolutional neural network for object detection, to first identify images containing people (Redmon and Farhadi, 2018). Then, using the bounding boxes provided by YOLO, we classified an image as a selfie if the primary subject appeared sufficiently close to the image border, a heuristic used to approximate arm's length photography.

We tested three bounding box proximity thresholds (99.9%, 99%, and 95% of the image dimensions). For each threshold, we further distinguished between individual and group selfies based on the number of people detected in the image.

To evaluate model performance, we manually labeled a stratified sample of 1,830 images, classifying each as an individual selfie, group selfie, or non-selfie. We selected the approach that yielded the highest weighted F1 score.

The selected model achieved a weighted F1 score of 0.75 for detecting individual selfies and 0.98 for group selfies. While these scores reflect the presence of classification noise at the image level, we use this variable only at the aggregate (monthly-city) level, where such noise is attenuated.

Moreover, we validated the selfie-based measure of narcissism through a separate survey study. The proportion of individual selfies posted was positively correlated with participants' NPI-40 narcissism scores (see Appendix A1), supporting the variable's use as a behavioral proxy for narcissistic self-presentation.

3.3.2 Entrepreneurial Quantity

To examine the volume of entrepreneurial activity across geographic areas, we relied on Crunchbase data, which includes detailed records of startup founding events. Our measure of entrepreneurial quantity is defined as the number of startups founded in a

given city during a specific month. This variable serves as a proxy for entrepreneurial liveliness and captures temporal and spatial variation in startup formation. The dataset includes monthly observations from 2015 to 2018 for 33 major U.S. cities, matched to social media data collected from Instagram.

Instagram data were scraped following the platform's public API limits, yielding a large corpus of images geolocated to each of the cities under study. A deep learning model based on the YOLO architecture was applied to classify images as selfies. We define a selfie as an image where a human face is both detected and positioned near the edges of the frame, an indication that the person depicted is also the one taking the picture. This measure is referred to as the number of YOLO-detected selfies. In addition, we extracted the number of group selfies (images with more than one face), and the number of images containing the hashtag `#selfie`. These variables together capture different dimensions of selfie-based self-presentation behavior.

The monthly Instagram data were matched with Crunchbase startup data by city and period, resulting in a panel dataset at the city-month level. This panel structure enables the use of fixed-effects Poisson regression to assess the relationship between selfie prevalence and startup quantity, controlling for time-invariant city characteristics and common temporal shocks.

3.3.3 Entrepreneurial Quality

To assess the quality of entrepreneurial activity, we leveraged the Entrepreneurial Quality Index (EQI) provided by the Startup Cartography Project (SCP) (Andrews et al., 2022). The EQI is a predictive measure of startup quality based on multiple attributes, including the backgrounds of founders, the industry entered, and location-specific entrepreneurial ecosystems. Higher EQI values are associated with startups that are more likely to experience rapid growth or successful exits. The SCP dataset is available at the yearly level and covers a subset of the original 33 cities for the years 2015 and 2016.

Given this constraint, we aggregated the Instagram and Crunchbase data to the city-year level to match the temporal resolution of the EQI data. The datasets were merged using ZIP code-level mappings to ensure precise geographic alignment. The resulting dataset allows us to analyze the relationship between selfie-based behavioral proxies for narcissism and the average quality of startups in each city.

The EQI serves as the dependent variable in a panel fixed-effects regression model that accounts for unobserved heterogeneity across cities. This model allows us to isolate the role of narcissistic self-presentation in predicting the quality, not just the quantity, of entrepreneurial activity.

3.3.4 Control Variables

In both the quantity and quality regressions, we included a consistent set of control variables to account for confounding factors. The number of inhabitants was included to account for the fact that cities with larger populations are likely to exhibit greater entrepreneurial and social media activity by virtue of their size. The number of tourists was also controlled for, as tourist volume may influence the number of selfies posted—particularly in visually iconic or popular locations—thereby potentially inflating our narcissism proxies. To capture the overall level of social media activity in each city, we included the total number of images posted on Instagram. This variable helps to normalize selfie prevalence against the general posting behavior of users in a given location. Additionally, we controlled for the number of hashtag selfies, measured as the number of images tagged with `#selfie`, in order to account for broader semantic trends in selfie posting that may not reflect self-focused visual framing.

In the quality regressions, we also included the number of startups in each city as a covariate. This was necessary to isolate the effect of narcissistic behavior on entrepreneurial quality from the general relationship between entrepreneurial quantity and quality. By controlling for startup volume, we aimed to assess whether narcissistic self-presentation

predicts the launch of high-quality ventures, independently of how many ventures are founded overall.

3.3.5 Estimation Strategy

To test the relationship between selfie prevalence and entrepreneurial activity, we employ panel data regression models that account for unobserved heterogeneity across cities. For the analysis of startup quantity, we estimate a fixed-effects Poisson regression model. This approach is appropriate for modeling count data with overdispersion and city-level fixed effects, and is robust to distributional assumptions under the quasi-maximum likelihood framework (Wooldridge, 1999). By including city fixed effects and time dummies, we control for both time-invariant location-specific characteristics and common temporal shocks.

For the analysis of startup quality, we use a fixed-effects linear regression model with the Entrepreneurial Quality Index (EQI) as the dependent variable. Given that EQI is continuous and approximately normally distributed, the linear specification allows for straightforward interpretation of coefficients. As in the quantity model, we include fixed effects to account for unobserved city-level differences and year-specific effects. All models use robust standard errors clustered at the city level to account for potential within-cluster correlation over time.

This dual-model approach enables us to distinguish the influence of narcissistic self-presentation—proxied by selfie prevalence—on both the frequency and the quality of entrepreneurial outcomes across geographic regions. Formally, our empirical models are specified as follows:

Startup Quantity (Model 1):

$$\mathrm{E}[Startups_{it} \mid X_{it}, \alpha_i, \gamma_t] = \exp(\beta_1 \cdot \mathrm{Selfies}_{it} + \beta_2 X_{it} + \alpha_i + \gamma_t)$$

Startup Quality (Model 2):

$$EQI_{it} = \beta_1 \cdot \text{Selfies}_{it} + \beta_2 X_{it} + \alpha_i + \gamma_t + \epsilon_{it}$$

where i indexes cities, t indexes time (month or year), α_i are city fixed effects, γ_t are time fixed effects, and X_{it} is the vector of control variables.

3.4 Results

3.4.1 Model 1: Entrepreneurial Quantity

Variables	**Entrepreneurial Quantity**
YOLO-detected selfies	0.000497***
	(0.000112)
YOLO-detected Group selfies	0.000176
	(0.000127)
Hashtag selfies	0.000142
	(0.000120)
Total images	-1.03e-05***
	(2.75e-06)
Population	2.52e-07
	(1.99e-07)
Tourist count	1.64e-09
	(3.64e-09)
Observations	1,343
Number of towns	32

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Table 3.1: Effect of Selfie Prevalence on Startup Quantity (Poisson Fixed Effects)

Table 3.1 presents the results of the fixed-effects Poisson regression. We find a significant and positive association between the number of YOLO-detected selfies and the number of startups founded per city-month. The coefficient on YOLO-detected selfies is statistically significant at the 1% level ($\beta = 0.000497$, $p < 0.01$), suggesting that higher prevalence of individual selfies is associated with increased entrepreneurial activity. Group selfies and hashtag selfies show no significant association.

This result indicates that narcissistic self-presentation behaviors are correlated with greater entrepreneurial dynamism at the regional level.

3.4.2 Model 2: Entrepreneurial Quality

Variables	**Entrepreneurial Quality (EQI)**
YOLO-detected selfies	-4.65e-08
	(1.67e-07)
YOLO-detected Group selfies	6.17e-08
	(5.66e-08)
Hashtag selfies	1.08e-07
	(9.67e-08)
Total images	-2.43e-09
	(2.84e-09)
Population	5.58e-09**
	(2.49e-09)
Tourist count	0.00
	(5.06e-11)
Startup quantity	1.15e-06**
	(5.25e-07)
Observations	444
Number of towns	21
R-squared (within)	0.336

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Table 3.2: Effect of Selfie Prevalence on Startup Quality (Fixed Effects)

Table 3.2 reports the results of the fixed-effects linear regression. The coefficient on YOLO-detected selfies is negative but not statistically significant ($\beta = -4.65 \times 10^{-8}$, $p = 0.357$), indicating no evidence that narcissistic behavior correlates with higher startup quality. Control variables perform as expected: population and startup count are positively associated with EQI.

Taken together, the results suggest that while narcissism may influence the quantity of startups, it does not appear to affect their expected quality or growth potential.

3.5 Discussion and Conclusion

This study explores how narcissistic self-presentation, proxied by selfie prevalence on Instagram, relates to entrepreneurial activity and startup quality at the city level.

Drawing on over five million geolocated images across 33 cities, we apply deep learning methods to detect selfies and combine them with data on startup quantity and quality. Our analysis reveals a nuanced relationship: while selfie prevalence is positively associated with the number of startups founded in a city, it does not predict the quality of those startups.

These findings contribute to the growing literature on narcissism and entrepreneurship in several ways. First, we extend prior individual-level work on narcissism and entrepreneurial intentions (Yadav and Batra, 2023) by showing that narcissistic behavior—manifested in publicly visible digital expression—also scales to aggregate entrepreneurial outcomes. The positive association between narcissistic self-presentation and startup quantity supports the view that narcissistic traits, such as confidence and risk tolerance, can stimulate entrepreneurial action. This is consistent with existing theories that link narcissism to a higher likelihood of venture creation due to enhanced perceived behavioral control and need for status (Navis and Ozbek, 2017; Engelen et al., 2016).

However, our results also show that narcissistic tendencies do not appear to enhance the quality of new ventures. This aligns with earlier research suggesting that narcissism may fuel bold decisions and innovation but also undermine long-term effectiveness due to overconfidence, impulsivity, or poor team management (Chatterjee and Hambrick, 2007; de Vries and Miller, 1985). Our null results for startup quality indicate that while narcissism might increase the quantity of entrepreneurial attempts, it does not guarantee higher success potential.

Theoretically, this paper adds to a growing stream of work using aggregated psychological characteristics to explain economic phenomena at the regional level. Prior research has shown that personality traits cluster geographically and shape entrepreneurship culture (Obschonka, 2017; Ebert et al., 2022). Our study builds on this literature by operationalizing narcissism through publicly observable behaviors, rather than self-report measures, offering a scalable and externally valid method to measure psychological traits

across cities. By leveraging computer vision to classify selfies, we provide a novel application of deep learning in the behavioral sciences that can be extended to other traits and outcomes.

Practically, these results offer insights for policymakers and entrepreneurship stakeholders. Regions with high levels of narcissistic self-presentation may exhibit vibrant startup ecosystems but could benefit from support mechanisms that promote strategic planning, team cohesion, and sustainable growth. Encouragingly, digital traces like selfies can serve as early indicators of entrepreneurial energy within urban populations.

Naturally, the study has limitations. First, while our proxy for narcissism is grounded in prior research, selfies may also reflect cultural or aesthetic trends unrelated to personality traits. Second, the quality measure (EQI) is only available for a limited period and subset of cities, restricting the generalizability of findings. Future work could expand the time window or use alternative quality indicators such as funding or exit success.

In sum, our findings show that narcissistic self-presentation may help spark entrepreneurial activity but does not predict higher-quality outcomes. As such, narcissism may be a catalyst—but not a guarantee—for entrepreneurial success. This distinction underscores the importance of evaluating both the quantity and quality of entrepreneurial activity when assessing the economic consequences of psychological traits.

A1 Appendix: Survey Validation of the Narcissism Proxy

To validate our behavioral proxy for narcissism, we conducted a supplementary survey in which participants completed the full Narcissistic Personality Inventory (NPI-40) (Raskin and Terry, 1988). Participants were asked to indicate their Instagram handle, and for those with publicly accessible profiles, one of the authors manually coded each account. Specifically, we recorded the number of total posts, the number of group selfies, and the number of individual selfies. This allowed us to compute the share of individual selfies posted by each participant.

The survey was administered through an internal university platform. All respondents provided informed consent, and the study was approved by the university's ethics committee. Participants received compensation in proportion to the time required to complete the questionnaire.

To assess the relationship between narcissistic traits and selfie behavior, we regressed the average NPI-40 score on the proportion of individual selfies posted. The results, shown in Table A1, indicate a significant positive relationship ($\beta = 0.30$, $p = 0.049$), suggesting that individuals who share more individual selfies tend to report higher levels of narcissism.

	(1) NPI-40 Score
Share individual selfies	0.303**
	(2.06)
Constant	1.326***
	(19.71)
Observations	29

t statistics in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table A1: Effect of Individual Selfies on NPI-40 Score

To better visualize how selfie-related variables correlate with each other and with narcissism, we report pairwise Pearson correlations in Table A2. As noted, the NPI-40 score is significantly positively correlated with the number of individual selfies ($r = 0.37$, $p = 0.048$).

	NPI-40	**Num. Posts**	**Num. Selfies**	**Indiv. Selfies**	**Group Selfies**	**Prop. Indiv.**	**Prop. Group**
NPI-40	1.0000						
Num. Posts	0.208	1.0000					
	(0.280)						
Num. Selfies	0.260	0.747***	1.0000				
	(0.173)	(0.000)					
Indiv. Selfies	0.370*	0.793***	0.944***	1.0000			
	(0.048)	(0.000)	(0.000)				
Group Selfies	0.189	0.765***	0.669***	0.728***	1.0000		
	(0.326)	(0.000)	(0.0001)	(0.000)			
Prop. Indiv.	0.369*	-0.090	0.165	0.256	-0.025	1.0000	
	(0.049)	(0.643)	(0.391)	(0.180)	(0.896)		
Prop. Group	-0.202	-0.038	0.032	0.009	0.298	-0.173	1.0000
	(0.293)	(0.843)	(0.868)	(0.964)	(0.117)	(0.369)	

Table A2: Pairwise Correlations: NPI-40 and Selfie Behavior

This study lends further support to our interpretation of individual selfies as a behavioral expression of narcissism.

References

Adjerid, I. and Kelley, K. (2018). Big data in psychology: A framework for research advancement. *American Psychologist*, 73(7):899.

Andrews, R. J., Fazio, C., Guzman, J., Liu, Y., and Stern, S. (2022). Reprint of "the startup cartography project: measuring and mapping entrepreneurial ecosystems". *Research Policy*, 51(9):104581.

Barry, C. T., Doucette, H., Loflin, D. C., Rivera-Hudson, N., and Herrington, L. L. (2017). "let me take a selfie": Associations between self-photography, narcissism, and self-esteem. *Psychology of popular media culture*, 6(1):48.

Belk, R. W. (2013). Extended self in a digital world. *Journal of consumer research*, 40(3):477–500.

Brandstätter, H. (2011). Personality aspects of entrepreneurship: A look at five meta-analyses. *Personality and individual differences*, 51(3):222–230.

Campbell, W. K. and Foster, J. D. (2011). The narcissistic self: Background, an extended agency model, and ongoing controversies. In *The self*, pages 115–138. Psychology Press.

Campbell, W. K., Goodie, A. S., and Foster, J. D. (2004). Narcissism, confidence, and risk attitude. *Journal of behavioral decision making*, 17(4):297–311.

Chatterjee, A. and Hambrick, D. C. (2007). It's all about me: Narcissistic chief executive officers and their effects on company strategy and performance. *Administrative science quarterly*, 52(3):351–386.

de Vries, M. F. K. and Miller, D. (1985). Narcissism and leadership: An object relations perspective. *Human Relations*, 38(6):583–601.

Dixon, S. (2023). Global social networks ranked by number of users. Statista.

Ebert, T., Gebauer, J. E., Brenner, T., Bleidorn, W., Gosling, S. D., Potter, J., and Rentfrow, P. J. (2022). Are regional differences in psychological characteristics and their correlates robust? applying spatial-analysis techniques to examine regional variation in personality. *Perspectives on Psychological Science*, 17(2):407–441.

Ellis, H. (1898). Auto-erotism: A psychological study. *Alienist and Neurologist (1880-1920)*, 19(2):260.

Engelen, A., Neumann, C., and Schmidt, S. (2016). Should entrepreneurially oriented firms have narcissistic ceos? *Journal of Management*, 42(3):698–721.

Fox, J. and Rooney, M. C. (2015). The dark triad and trait self-objectification as predictors of men's use and self-presentation behaviors on social networking sites. *Personality and Individual Differences*, 76:161–165.

Gebru, T., Krause, J., Wang, Y., Chen, D., Deng, J., Aiden, E. L., and Fei-Fei, L. (2017). Using deep learning and google street view to estimate the demographic makeup of neighborhoods across the united states. *Proceedings of the National Academy of Sciences*, 114(50):13108–13113.

Leung, Y. K., Franken, I., Thurik, R., Driessen, M., Kamei, K., Torres, O., and Verheul, I. (2021). Narcissism and entrepreneurship: Evidence from six datasets. *Journal of Business Venturing Insights*, 15:e00216.

Navis, C. and Ozbek, O. V. (2017). Why context matters: Overconfidence, narcissism, and the role of objective uncertainty in entrepreneurship. *Academy of Management Review*, 42(1):148–153.

Obschonka, M. (2017). The quest for the entrepreneurial culture: psychological big data in entrepreneurship research. *Current opinion in behavioral sciences*, 18:69–74.

Obschonka, M., Lee, N., Rodríguez-Pose, A., Eichstaedt, J. C., and Ebert, T. (2020). Big data methods, social media, and the psychology of entrepreneurial regions: capturing cross-county personality traits and their impact on entrepreneurship in the usa. *Small Business Economics*, 55:567–588.

Raskin, R. and Terry, H. (1988). A principal-components analysis of the narcissistic personality inventory and further evidence of its construct validity. *Journal of personality and social psychology*, 54(5):890.

Redmon, J. and Farhadi, A. (2018). Yolov3: An incremental improvement. *arXiv preprint arXiv:1804.02767*.

Sorokowski, P., Sorokowska, A., Oleszkiewicz, A., Frackowiak, T., Huk, A., and Pisanski, K. (2015). Selfie posting behaviors are associated with narcissism among men. *Personality and individual Differences*, 85:123–127.

Weiser, E. B. (2015). # me: Narcissism and its facets as predictors of selfie-posting frequency. *Personality and individual differences*, 86:477–481.

Wooldridge, J. M. (1999). Distribution-free estimation of some nonlinear panel data models. *Journal of Econometrics*, 90(1):77–97.

Yadav, R. and Batra, S. (2023). Does narcissism influence entrepreneurial intentions? a theory of planned behaviour perspective. *The Journal of Entrepreneurship*, 32(2):449–478.