In space, there is no-one to influence: An enquiry into how and why Instagram influencers are made by their readers.

Alexander Nicholas Shepherd

Fourth Year Project Report
School of Informatics
University of Edinburgh
2020

Abstract

This report aims to answer how Instagram influencers are influential by assessing Influence Space, which compromises of 3 dimensions (Authority, Persuasion and Topic Initiation and Uptake) and 3 levels (Global, Domain and Hashtag Topic). From the functions, which we have developed in synergy with literature spanning over sociology through to computer science, we found that the extent of authority and persuasion display collaboration between influencers and their readers.

In the authority dimension of Influence Space, we found that among influencers, high authority is associated with more readers adopting an idea a small number of times, with influencers posting the minimal amount in a hashtag conversation. In the persuasion dimension, we found that influencers tend to show a community effect, where they achieve homophily with readers rather than "persuade" readers. However, when persuasion occurs, then readers make most of the effort in the interactions. In the topic initiation and uptake dimension, we found that the nature of the influencer's domain determines how likely an influencer is to set a trend.

Table of Contents

Chapter 1

Introduction

Following the birth and massive adoption of social media, we see a new form of ”celebrity” appearing on social media platforms, which are known as social media influencers. The question that is asked by many social media users is: How did these users become influencers?

This report aims to provide an insight into this question and will argue that influencers are primarily influential because of the reader’s behaviour and not necessarily of the influencer.

This similar line of thinking is also supported in Watts and Dodds, 2007, who believe that influencers are ”accidental influentials”: their status comes about as an ”accident of location and timing”. in other words, their status is a product of their audience and not of influencers’ behaviours.

So far, the problem with research into influence is two-fold:

Firstly, it is currently somewhat dichotic. On one side, the literature characterises influence naively as a single behaviour, which is an oversimplified evaluation of influence, as per the definition from Biran et al., 2012. On the other side, research into the specific influential behaviours is particularly sporadic and lacks a unified focus around influence, particularly in relation to social media influencers.

Secondly, despite multiple theories covering influential behaviours arguing that influential behaviours are group processes (e.g. Platow et al., 2015), most research counter-intuitively avoid looking into an influencer’s audience when assessing an influencer’s degree of influence.

Therefore, this report takes a unique approach to defining social media influence by quantifying specific influential behaviours, which make up ”Influence Space”. Here we will support our proposed measures with relevant theories and computational formulations from a variety of research fields, which do aim to measure readers' as well as influencers' activity. We will then gain an overall perspective on how an influencer is particularly influential by measuring and comparing these influential behaviours over 100 Instagram influencers across 5 lifestyle domains (Beauty, Fashion, Fitness, Food and Travel) with 100 "regular users", who are not influencers.

This report will provide valuable insight into fields of continuing research, such as influencer detection, which can help businesses to identify individuals who can represent their brands and increase sales. This research will also for all social media users generally, who want to understand how to gain greater influence in social media.

To achieve this research goal, we will describe and discuss the following in their respective sections:

- 2.1 Define an influencer and their potentially distinct influential behaviours, which correspond to the 3 dimensions of Influence Space. We will also define the 3 levels of influence, which focuses on the span of a user's influence: From global (across the social media platform), through domains to specific hashtag topics.

- 2.2 Establish a rounded and detailed understanding of each dimension of Influence Space from the economic, sociological and sociolinguistic literature.

- 3 Define the dataset used in the analysis, which comprises 33 million Instagram posts, 19 million users, in 1,000 hashtag conversations, over the period 2011 - 2019.

- 4.1 - 4.3 Develop functions to quantify each dimension of Influence Space, building on literature reviewed in Section 2.

- 4.4 Define hypotheses to understand and compare both influencer's and regular user's influential behaviour over all 3 dimensions in Influence Space, and across all 3 levels of influence.

- 5 Provide results of the analysis for each hypothesis, using functions and data defined in Sections 3 & 4 respectively.

- 6 Discuss the results to identify particularly notable influential behaviours, which distinguish influencers from regular users, across all levels of influence.

- 7 In conclusion, summarise the key findings which evidence that the source of influence is derived predominantly from the reader and not the influencer. We will also put forward ways that influencers can best harness the power of their readers to grow and sustain influence.

This project has produced a robust framework which helps to understand social media influencers' behaviour and the role of the reader in each of the influential behaviours. We have done this by developing functions, which focus on interactions between authors and readers that extend beyond the limitations of the passive influencer-follower relationship. The framework has been motivated by a broad spectrum of theories, across several research fields, from which we have built formulations to quantify influential behaviour.

Chapter 2

Background Research

2.1 Influencers in Influence Space

Viewing influence as a single behaviour is naive. This is illustrated by (Cha et al., 2010), which aims to identify influence among popular Twitter users through a single behaviour. This study, however, lacks an understanding of influence.

Firstly, the concept of influence is not explicitly defined, with reference to literature. Secondly, the study subconsciously focuses on authority, one of several influential behaviours.

Therefore, we need to view influence as a set of influential behaviours. Cialdini and Goldstein, 2004's highly-cited literature review on social influence, provides a necessary motivation to view influence over multiple dimensions. They break down social influence into compliance and conformity. To yield compliance over an audience, an user must initially have authority over them, whether that is based on expertise or where they sit in the social hierarchy. Conformity is denoted as a user's change in their behaviour, such that they align to others (Cialdini and Goldstein, 2004) to gain social approval. This could be either changing one's opinions or conforming to a trend.

The literature (Biran et al., 2012, Rosenthal and McKeown, 2016 Nguyen et al., 2014) depicts an influencer, through the following influential behaviours, where each will serve as a dimension of what we term as Influence Space (See Figure 2.1):

2.2.1 **Authority**: An influencer has explicit authority in a group.

2.2.2 **Persuasion**: Other users express alignment towards influencer's opinion.

2.2.3 **Topic Uptake and Initiation**: Influencers introduce new trends within topics, which are later followed by other users.

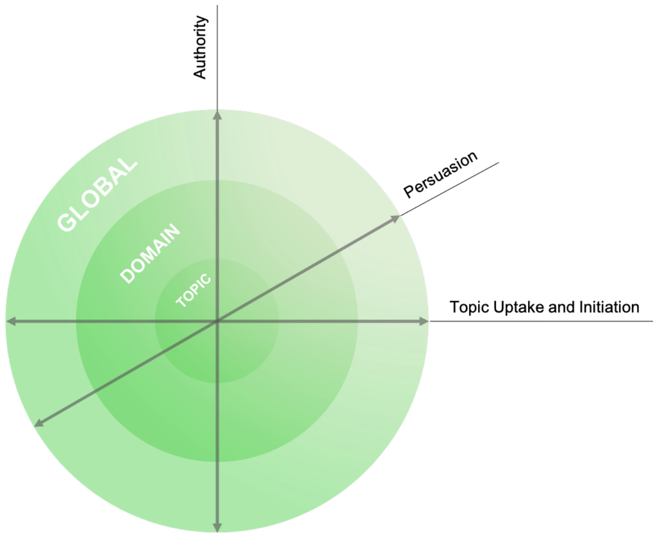

Figure 2.1: Visual depiction of the influence space, with the three dimensions of influence [Authority, Persuasion and Topic Initiation and Uptake] over the three levels [Global, Domain, Topic]

As stated earlier, the goal of the report is to develop an understanding of why and how influencers are influential. In our approach, we want to assess these dimensions in Influence Space over various levels. Such levels of influence are motivated by the fact that influencers are allegedly not globally influential (irrespective of domain or topic), but situationally influential (within domains and/or topics) Rosenthal and McKeown, 2016. Therefore, our analysis will review influential behaviour via the 3 dimensions in Influence Space over 3 levels of influence, which refer to span of a user's influence in a given dimension (See Figure 2.1).

2.2 Quantifying the Dimensions of Influence Space

In this section, we will use existing literature from a variety of research fields to provide an understanding of each dimension of Influence Space in the context of social media influencers. This is necessary to provide a solid foundation for the functions which we will later introduce to quantify each dimension (See Section 4).

To reiterate, the 3 influential behaviours, which correspond to the 3 dimensions of Influence Space, that we will explore in this paper are as follows:

- Authority
- Persuasion
- Topic Initiation and Uptake

2.2.1 Authority

The original definition of an influencer from Biran et al., 2012 defines an authoritative user as an opinion leader, whose ideas are adopted by others in their immediate environment Katz and Lazarsfeld, 1966.

In the literature, there are two main approaches from both social psychology and computer science, which quantify authority. One approach disregards Biran et al., 2012's definition and instead reviews static, user-based features that are attributed to the opinion leaders themselves (such as personality). The other approach aligns more to Biran et al., 2012's notion of authority being idea adoption by focusing on the group dynamics and connectivity of the leader in relation to other users.

In social psychology, several studies have unsuccessfully tried to show reasonable associations between five-factor personality traits (e.g. Neuroticism, Extroversion, etc.) and leadership success (Bartone et al., 2009, Judge et al., 2002), to which extroversion is seen as being most salient in leadership success. These studies involved directly testing and assessing individuals, which is not feasible in our analysis.

Alternatively, we can implicitly infer personality traits from one's use of language. For example, using Linguistic Inquiry and Word Count [LIWC] Pennebaker and King, 1999. Although these methods have been often used in the literature (e.g. Danescu-Niculescu-Mizil et al., 2011), the only identified paper, which assess the reliability of these methods, shows significant inconsistencies Flaxman and Kassam, n.d.

In addition, Platow et al., 2015 states that the significance of these traits can be explained by in-group similarity. In the case of extroversion, for example, a relatively large study Platow et al., 2006 shows that subjects ascribed greater charisma to individuals who were members of the same group than those outside the subjects group.

Platow et al., 2015 suggests that membership in a group conversation provides the environment in which an influencer can be an opinion leader at a given point in time. He continues to argue that authority is a group process, and is not something that an individual can possess on their own without an audience. You cannot have authority or be a leader, if you have no-one to follow your lead.

Watts and Dodds, 2007 builds on this idea and suggests that being a opinion leader is associated with the connectivity between individuals. This line of inquiry has been thoroughly explored in the literature.

An approach proposed by Cha et al., 2010 aims to associate influence with an individual's connectivity in a network. They define influence by the following implicit connectivity features: number of followers (in-degree), re-tweets and user mentions measured across 3 popular topics, over a given time period. These features logically relate to an influencer's popularity, virality of content and user engagement respectively. For each of the 6 million users involved in their analysis, the features were measured and ranked. The study shows mild correlation ($\rho \approx 0.5$) in connectivity feature rankings over topics. This suggests that those who are authoritative over a certain popular topic, can be relatively authoritative over other topics, and thus provides evidence for a global influencer authority. This study also demonstrates that, to measure authority, we also have to align closely to Biran et al., 2012's original definition, where authority is expressed as idea adoption.

The notion of idea adoption is shown to be closely associated with information diffusion in the literature (Beal and Bohlen, 1956, Bohlen, 1967, Valente and Davis, 1999).

Bakshy et al., 2011 explores idea adoption as an information cascade, by analysing continuous series of Twitter re-tweets that contain a certain link, which denotes a cascade. Influence scores are computed as the average log of the cascade size (number of users involved in cascade) from when the user was randomly-chosen to start the propagation of a given link (termed as "seed" user). This model effectively captures idea adoption globally over a network, by measuring the size of the information cascade. However, it randomly initiates the "seed" user, rather than attending to who first posted the link. Hence, it provides little insight about the relative authority of certain influencers. Furthermore, the study used a relatively short analysis window (two months between 13/09/19 - 15/11/19) and a limited number of topics, which raises questions as to how well the findings could be generalise over large-scale social media platforms like Instagram.

Myers et al., 2012 reviews connectivity with a greater level of sophistication. They propose that a user's propensity to adopt an idea is the sum of external and internal influence. Re-posting something from an influencer you follow on social media, for example, is deemed as internal influence. Posting something that is from a source external to who you follow or whom follows you is described as external influence. This line of thinking enables us to separate information diffusion as a result of authority versus independent posting events: many studies have failed to consider this important distinction.

Myers et al., 2012 propose that the probability for a user i to adopt an idea by both internal and external exposure to information can be modelled by the following pseudo-cumulative distribution function:

$$F^{(i)}(t) = \sum_{n=1}^{\infty} P(i \text{ has } n \text{ exposures }) \times P(i \text{ adopts idea } | i \text{ has } n \text{ exposures})$$

Here, idea adoption is simply modelled by multiply the probability that user i is externally exposed by the probability that user i is internally exposed. We also note that idea adoption is modelled over the number of times the user i is exposed to a certain piece of information. Although this formulation of idea adoption is intuitive, it requires complex optimisation of many unconstrained parameters, which influence the overall distribution function. The results from the paper focus on the efficacy of the formulation to identify volumes of external exposure over time rather than insights into idea adoption.

In summary, the literature shows authority measured in terms of personality traits and information diffusion. As we have mentioned before, measuring authority through influencer's and regular user's traits is simply too cumbersome to conduct. Even if we refer to methods, such as LIWC analysis, to implicitly extract these traits, evidence suggests that these forms of extracting traits, implicitly from language, is not robust.

We therefore revert to the original definition of authority from Biran et al., 2012, which views this dimension as idea adoption, measured by information diffusion models. In the literature, measuring idea adoption primarily relies on connectivity between users and their immediate environment, where the idea to be adopted is typically a URL link or a hashtag. Most of these formulations refer to implicit measures, similar to those used in social psychology (Bakshy et al., 2011), and more sophisticated, probabilistic formulations (Myers et al., 2012) require a number of working assumptions and output very little insight into idea adoption. A new probabilistic approach will need to be developed taking inspiration from the literature referenced above. The method we have developed is described in Section 4.1.

2.2.2 Persuasion

We define persuasion as the change in an individual's attitude through communication with others (Petty and Cacioppo, 2012, Wood, 2000). Miller, 2013 states that persuasion relies on the 'power of verbal [...] symbols' to incite a change in behaviour. The importance of communication in persuasion has also been highlighted in studies, including that from Coffman and Niehaus, 2014.

In the study, they measure persuasion by the consumer's change in valuation of a product before and after a communication condition. The communication condition involves a telephone conversation with a randomly-selected salesperson. In the non-communication condition, subjects waited on the call with no interaction for the same fixed length of time. The study revealed that the salesman communication showed a significant average increase in the consumer's valuation of the product. This study, however fails to relate the change in attitude as a function of the communication itself, which is prescribed in the definition of persuasion.

For the rest of this section, we will focus on theories and formulations which evaluate to what extent one is persuasive over communication with readers and how exactly does an author persuade a reader (the route of persuasion).

2.2.2.1 The extent of persuasion

The Communication Accommodation Theory (CAT) Shepard, 2001 explains the changes in attitude through an individual's use of language when communicating. In CAT, the individual's use of language (i.e. linguistic style) is used as a way to achieve a certain social distance (or closeness) between themselves and others. For example, readers can be socially closer to fashion influencers by simply including more vocabulary which the influencers use. This example reflects the notion of linguistic accommodation, in which the reader changes their linguistic style, influenced by the author's linguistic style. In effect, linguistic accommodation captures the central principle of persuasion.

The theory continues to talk of the factors to include when approximating linguistic accommodation between users:

- Degree of accommodation: e.g. how much does user A change their linguistic style?
- Direction of accommodation: e.g. does user A change their linguistic style towards or away the direction of user B? Do both users A & B change their respective linguistic styles, such that they are closer together (converge)?

Shepard, 2001 stresses that these factors, of course, can vary depending on the groups of users a particular user is communicating with. This is because set of users define a set of norms; shared ideas about what is appropriate and inappropriate when communicating within a certain group of users. For example, a subordinate would change their linguistic style to accommodate, for example, for speaking with a superior in a formal setting (Azuma, 1997). What is defined as an appropriate linguistic style with friends, does not necessarily mean that the same linguistic style is appropriate with a superior.

In this particular setting, the subordinate would presumably enforce a greater degree of accommodation towards the superior, for example, than the superior to the subordinate. This is an example of asymmetric accommodation, and is of course, very intuitively associated with persuasion. But take a different situation, where author a and reader r accommodate symmetrically towards each other. What does that mean?

CAT and Šćepanović et al., 2017 argues that this is an example of homophily; the idea that users are similar to each other (De Choudhury et al., 2010). Šćepanović et al., 2017 goes on to differentiate linguistic accommodation from homophily. Accommodation concerns asymmetric changes in linguistic style relative to another user, whilst homophily concerns symmetrical convergence in linguistic styles in response to each other's change in linguistic styles. For these reasons, we will argue, that both homophily (symmetrical) and asymmetric accommodation are considered to be extents of persuasion, i.e. how much a reader changes their attitude towards (is persuaded by) the author, relative to the author's own changes in attitude.

There has been many examples of computing the extent of persuasion, which is normally referred to in the literature as Stylistic Accommodation. Although the formulations themselves vary, we can split these formulations into how they capture linguistic style from text. One approach maps linguistic style into syntactic space and the other into semantic space. We will refer to mapping into syntactic space as simply counting the number of specific word classes (e.g. 1st person singular pronouns, quantifiers, etc.). We will refer to semantic space as interpreting linguistic style as a document embedding described further in *Viewing linguistic style in Semantic Space*.

In the following subsections, we will review and discuss both of these approaches that capture the extent of persuasion, relative to the definition of persuasion.

Viewing linguistic style in Syntactic Space

One of the first to propose a Stylistic Accommodation formulation, is work from Danescu-Niculescu-Mizil et al., 2011, which is allegedly motivated by CAT. This formulation works on the basis that linguistic style is derived from mapping text into syntactic space, via the Linguistic Inquiry and Word Count analysis (LIWC) Pennebaker and King, 1999. The motivation behind using syntactic space was to differentiate linguistic style from content. The count of each word class in syntactic space represents a style dimension.

The formulation is based on a probabilistic framework, whereby the stylistic accommodation for a style dimension C over texts T between author a and reader r was computed as the following:

$$Acc_{a,r}(C) = P(T_r^C | T_a^C, T_r \rightarrow T_a) - P(T_r^C | T_r \rightarrow T_a)$$

$T_r \rightarrow T_a$ refers to the text from reader r is a direct response to the text from user a.

This formulation simply computes the difference between probabilities between r's use of a style dimension, given that a has also used that style dimension and r is responding to a and the r's use of a style dimension given that they respond to a.

This formulation's use of a probabilistic framework, which deduces the probability of a particular style dimension will change given a user's use of that style dimension. This can separate instances of persuasion from instances where users have posted independently of each other. However, it fails to truly reflect the principles of CAT. Firstly, the probabilistic framework does not explicitly take into account both the degree and direction of accommodation which was explicitly defined in CAT. Secondly, this formulation solely focuses on stylistic accommodation being an indicator of persuasion and excludes measurement of homophily.

Jones et al., 2014 builds on Danescu-Niculescu-Mizil et al., 2011 by proposing the Zelig quotient, which identifies the change in author a's linguistic style from their baseline style with respect to an audience of readers R's linguistic style, which is represented in the exact same way in syntactic space as Danescu-Niculescu-Mizil et al., 2011.

They define user a's baseline style μ_a as the following:

$$\mu_f(a) = \sum_{m=1}^{n_a} \frac{f_m(a)}{n_a}$$

, where n_a represents the number of posts a has authored and $f_m(a)$ the value for a given a particular style dimension in syntactic space $f \in F$ for a given post m.

Before computing the quotient, they firstly compute the degree of accommodation from the author a to the reader r as the following:

$$f(a,r) = \sum_{m=1}^{n_a r} \frac{f_m(a,r)}{n_a r}$$

, where $n_a r$ represents the number of posts user a has posted to the reader r.

The Zelig Quotient for user a $Zelig(a)$ is computed by the following:

$$Zelig(a) = \frac{1}{R_a} \sum_{r=1}^{R_a} 1 - (\frac{|\vec{ar}|}{|\vec{\mu r}|})(\frac{|\vec{\mu r}|^2 + |\vec{\mu r}|^2 - |\vec{ar}|^2}{2(\vec{\mu r} \cdot \vec{ar})})$$

, where R_a represent the number of users, which respond to the author a's post, $a = f(a,r) \in F$.

Here a quotient of 1 indicates that the author a is readily accommodating to the audience of readers $r \in R$. A negative quotient indicates divergence between the author and their audience.

From an initial analysis over three online communities, they found that non-leaders (i.e. regular users) show a significantly higher Zelig quotient on average than leaders (i.e. influencers). Therefore, we expect that regular users are typically readily accommodate to their audience than influencers.

This formulation does adopt more of the principles of CAT; for example, it assesses the change in author's linguistic style, relative to the audience. It also enables us to assess the degree of accommodation by comparing the author's linguistic style from one posts, against their overall baseline (typical) style.

However, like with Danescu-Niculescu-Mizil et al., 2011's formulation, it overlooks the change in the reader's individual linguistic style (i.e. comparing readers' baseline styles with their response to the author's post) to account for homophily as a form of persuasion.

Furthermore, in both formulations, there is a limitation in utilising syntactic space to represent linguistic style. It does not account for changes in semantic meaning, which is also involved in stylistic accommodation Reitter, 2017. For example, if the author posts ”I like eating a lot of vegetables.” and a reader posts ”I like eating lots o’ cake.” Firstly, given the more casual use of language in social media (e.g. contractions [e.g. ”o’” instead of ”of”] and misspellings), counts over word classes will be underrepresented and hence could falsely predict that the above messages have similar linguistic styles. Secondly, and more importantly, in this example, the syntactic space does not clearly capture changes in attitudes, which is an essential premise of persuasion. In the example posts, the author demonstrates a positive attitude towards vegetables and the other cake, which in the Zelig quotient’s formulation, would paradoxically interpret these posts as similar, given they have similar counts in word classes.

Therefore, to measure persuasion according to our specified definition, we will represent linguistic style in semantic space.

Viewing linguistic style in Semantic Space

The only known formulation of Stylistic Accommodation to utilise linguistic styles in semantic space is that from Nasir et al., 2019. In their study, they view stylistic accommodation as computing semantic cohesion between the author a and reader r as a Word Mover’s Distance (WMD) between their linguistic styles.

Originally proposed by Kusner et al., 2015, the WMD is used to translate a set of word embeddings involved in a document into distances between documents. The method proposes the use of word2vec models to produce word embeddings (for an explanation of how word/doc2vec models work, see Section 4.2). This method of producing word embeddings has been shown to capture both semantic and syntactic information (Mikolov et al., 2013).

For a pair of word embedding vectors from post a and r (w_a and w_r respectively), we compute the distance between them as the following:

$$dist(w_a, w_r) = ||\sqrt{\sum_{d=1}^{D} w_{ad}} - \sqrt{\sum_{d=1}^{D} w_{rd}}||$$

, where $\sqrt{\sum_{d=1}^{D} w_{ad}}$ represents the Euclidean distance for the word embedding w_a.

The WMD between two posts a and b is then computed by the following:

$$WMD(a,b) = \min_{T \geqslant 0} \sum_{i=1}^{m} \sum_{j=1}^{n} T_{ij} dist(w_i, w_j)$$

, subject to

$$\sum_{j=1}^{n} T_{ij} = \frac{c_i^1}{n} \quad \forall_i \in \{1, ..., m\}$$

, and

$$\sum_{i=1}^{m} T_{ij} = \frac{c_j^2}{m} \quad \forall_j \in \{1, ..., n\}$$

, where m, n represent the number of unique words in posts a and r respectively, c_i^k represents the frequency of word w_i in post k, and T is a matrix which can be interpreted as finding associations between neighbouring words in semantic space and can be optimised using approximation techniques.

From the WMD between posts, Nasir et al., 2019 formulate stylistic accommodation between users a and r as the following:

$$Acc(a, r) = \frac{1}{N} \sum_{i=1}^{N} dist_i^{A \rightarrow R}$$

, where

$$dist_i^{A \rightarrow R} = \min_{i \le j \le i+k-1 \le N} WMD(A_i, R_j)$$

, where N represents the number of posts

This equation essentially finds the average minimum WMD in a conversation between an equal number of posts from the users a and r, where the author a posts before reader r.

The formulation does align with many of the principles of CAT. WMD captures the degree of accommodation. Furthermore, mapping linguistic style into semantic space provides information about a user's attitude. From this, we can develop a formulation that aligns with the definition of persuasion.

The formulation, however, still fails to explicitly account for homophily as a form of persuasion. There is no output value for the function $Acc(a, r)$, which corresponds to symmetrical convergence. Furthermore, there is no explanation as to how WMD is computed between posts of varying lengths (number of words in a post), which will most likely to be the case in a dataset of social media posts.

2.2.2.2 The routes to persuasion

Based on sociological research into persuasion, the Elaboration Likelihood Model (ELM) Petty and Cacioppo, 2012 aims to explain how attitudes change as a function of communication.

ELM proposes two routes to persuasion: one which requires reasoning (i.e. a logical argument) based on provided evidence, defined as the central route and the other relies on superficial, simpler cues (e.g. attractiveness of messenger) defined as the peripheral route.

Changes in attitudes from the central route are more likely to remain stable over time, whilst changes in attitudes from the peripheral route are more likely to fluctuate over time. For example, when we are purchasing an item, if we actively review information about the item and see well-argued recommendations (e.g. from influencers), then we are likely to develop a stable attitude towards the item. However, if an influencer says "buy this item" without any elaboration, then your attitude towards the item is likely to fluctuate over time.

In summary, we understand persuasion to be a change in attitudes in response to another's change in attitudes. From CAT (Shepard, 2001), we need to consider both the degree and direction of accommodation. We also framed a user's attitude as a linguistic style. Users who change their linguistic styles equally towards each other is deemed to homophily (Shepard, 2001, Šćepanović et al., 2017). We have also argued that as homophily looks at a change in linguistic style with respect to another user's change in linguistic style, then we should also consider homophily (symmetrical convergence) to be, by definition, a form of persuasion.

Varying formulations have tried to encapsulate the idea of Stylistic Accommodation, motivated by CAT. Some have formulated Stylistic Accommodation by representing linguistic styles over syntactic space, which is argued to be ineffective. Representing linguistic styles has been argued to be beneficial in corresponding to changes in attitude. However, the current, proposed formulations still falls short of considering homophily into the accommodation formula. Therefore, we need to account for accommodation symmetry, which will help to discriminate homophily from asymmetric accommodation.

Finally, we looked at the routes to persuasion, which aim to explain how an influencer persuades their audience. The central route to persuasion denotes using logic and reasoning to persuade readers and is shown by stable linguistic style over time after the reader has been persuaded. The peripheral route to persuasion utilises superficial features, such as the influencers appearance or personality and is shown as by variable linguistic style over time after the reader has been persuaded.

2.2.3 Topic Initiation and Uptake

Biran et al., 2012 defines topic initiation and uptake as a way to measure how influencers respond to and participate in user's interests and trends. However, from this definition, Biran et al., 2012 takes a different approach. They interpret this as demonstrating the influencer's ability to change the course of the conversation via dialog patterns. This interpretation however proves to have very mixed results in identifying influencers. Furthermore, the methodology behind detecting these dialog patterns is not noted in the paper, and therefore we cannot replicate findings.

After conducting a literature search for trend setting in social media, it is clear that very little research has been published in this space. The only example eluding to identifying trends is from Shen et al., 2017.

The goal of this paper is to identify trends (as well as influencers) on Twitter, by reviewing the descriptive statistics (mean, standard deviation) over the number of daily and monthly posts in separate conversations about ”mobile learning” and ”online learning” over time. This method simply tracks how these conversations change over time. Each conversation involve posts with certain hashtags in them.

Although this formulation tracks conversation changes by number of posts per unit time, it does not take account of the author's contribution to that conversation. This is vital in understanding how exactly an influencer initiates or responds to trends.

Chapter 3

Data

Although, we have seen a promising surge in the number of publications regarding influence in social media platforms, particularly Twitter (e.g. Danescu-Niculescu-Mizil et al., 2011, Cha et al., 2010), to-date there is no publication which assess influence in the context of Instagram; a platform which shows noticeably higher activity than Twitter (See Figure 3.1).

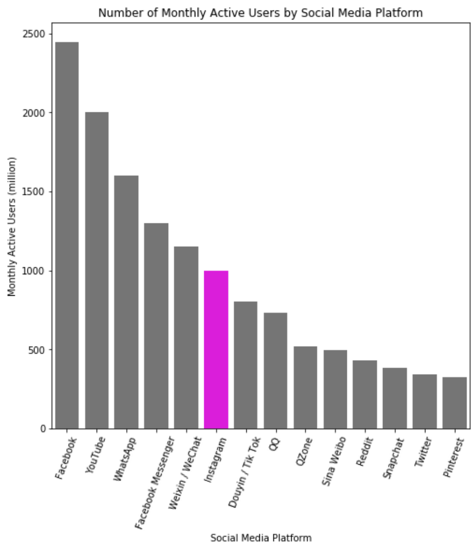

Figure 3.1: Number of monthly active users by social media platform as of January 2020. Data taken from "Global social media ranking 2019", 2020

Since there were no studies of influence among Instagram influencers, we have had to design our own approach to identify the top 100 Instagram influencers and 100 regular Instagram users. The influencers are equally split over 5 lifestyle domains, namely: Beauty, Fashion, Travel, Food and Fitness. These influencer have been identified and validated from various online sources ("25 Top Foodie Influencers", 2017, "The Top

Fitness Influencers to Guide Your Workout", 2019, "Beauty Influencers: The Top 25 You Need to Follows", 2017, Agrawal, 2019, Forsey, 2019, "THE 10 TOP TRAVEL INSTAGRAM INFLUENCERS WHOSE JOURNEYS YOU SHOULD FOLLOW IN 2019", 2019), which have rated influencers by most influential in each of the domains.

In our analysis, we want to distinguish influential behaviours which are synonymous with influencers. To do this, we need to compare the behaviours of influencers from regular users. Therefore, for this analysis, we identified 100 randomly chosen Instagram users, which were either:

1. The influencer's number of followees (number of users influencers follow) : followers ratio is below the top 95th percentile of the ratio among the identified influencers.

2. The user's number of followers is below the top 99th percentile of number of followers among the identified influencers.

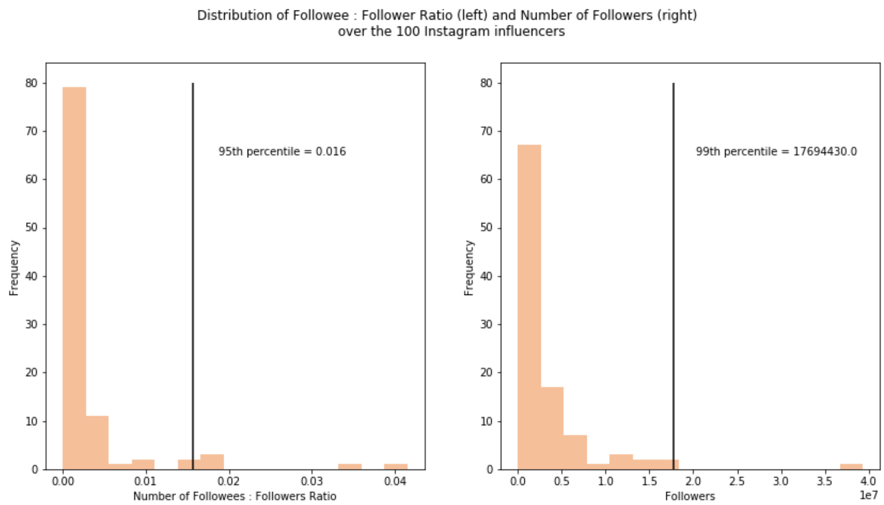

Figure 3.2: Distribution of Followee : Follower Ratio (left) and Number of Followers (right) over the 100 Instagram influencers. Regular users, which were randomly selected, were either: below 95th percentile over Followee : Follower Ratio or below the top 99th percentile over number of Followers

This selection criteria was motivated from the literature, which have used similar measures to identify influencers (Agam, 2017, Anger and Kittl, 2011, Sun and Ng, 2012, to identify a few).

We will use hashtag conversations to enable us to assess these dimensions in influence space.

We define a hashtag conversation simply by the Instagram posts which involve a specific hashtag topic. So a hashtag conversation normally has a large number of participants including the influencer and other users.

As mentioned in Section 2, to quantify the level of authority in the form of idea adoption, we need to identify a conversation over which authority can be derived. To quantify persuasion, we need to observe users' changing linguistic styles over a conversation. To measure topic initiation and uptake, we need to again observe these over a conversation.

We extracted hashtag conversations in the following manner:

Firstly, we performed a count of hashtags used by our 100 target influencers. From those counts, we computed the top 10th percentile over hashtag counts for each influencer. We then only kept those hashtag topics which exceeded the top 10th percentile for that particular influencer. This lead to a total of 7,432 hashtag topics.

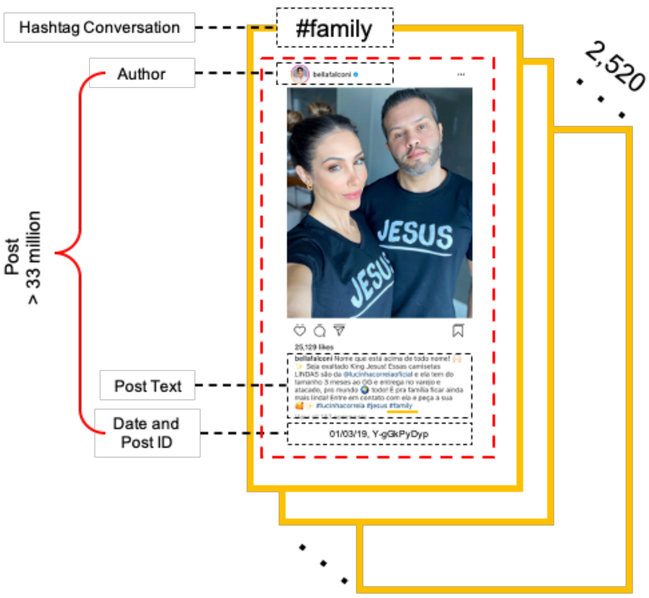

Figure 3.3: Visual representation of data extracted from Instagram, which is grouped into hashtag conversations.

Before extracting hashtag conversations, we found when any of the influencers first posted in that hashtag conversation and used this to prevent further unnecessary data extraction.

Thus the timeframe for the overall analysis spans over 8 years between 28/04/2011 to 8/11/2019.

Out of the 7,432, we have successfully extracted 2,520 hashtag conversations (due to time restrictions and the Instagram API extraction protocol, which fails to extract all posts from a given hashtag conversation.) which comprises of over 33 million posts and includes approximately 19 million users, including our target influencers and regular users.

Chapter 4

Methodology

In this section, we define the functions that will be implemented to quantify Authority, Persuasion and Topic Uptake/Initiation.

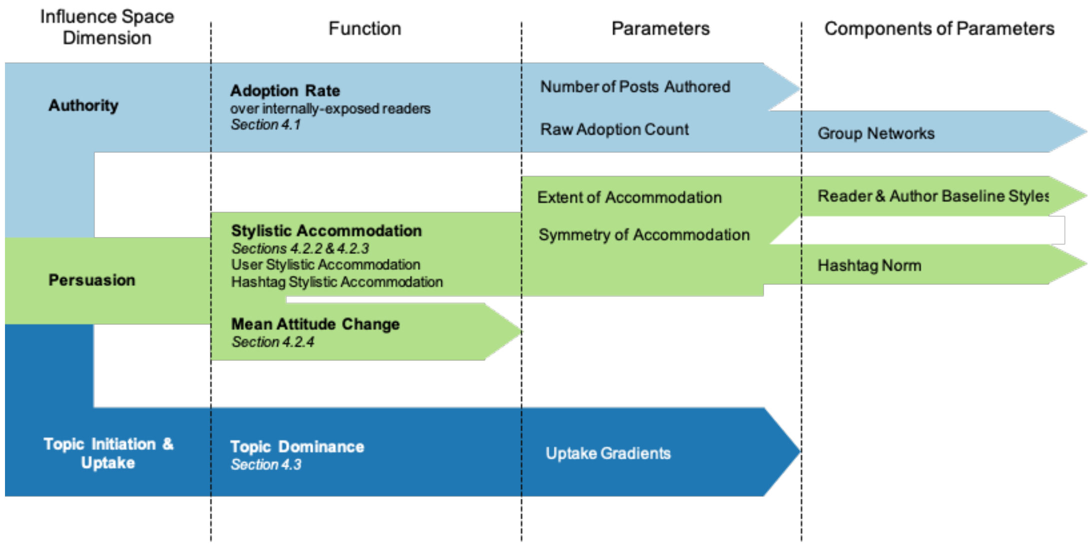

Figure 4.1: Overview of Methodology, which introduce the functions which will quantify their respective dimension of Influence Space

We also use the terms author and reader users to denote the post creator (author) and those that subsequently take part in the conversation (reader), and are influenced as a result. We need to stress that these are independent of user type (influencers and regular users).

4.1 Idea Adoption

From Section 2.2.1, we define authority as the level of idea adoption among readers over a whole conversation. Furthermore, we also mentioned that idea adoption is synonymous with information diffusion and so we must consider connectivity between users. In this section, we will explain how we quantify idea adoption and the assumptions we made when defining the functions.

The typical way to consider connectivity between users is to explicitly understand which users follow other users in the network. However, the scale of network renders it impractical to extract the network from the Instagram API within the timescale of the report. For example, from Figure 4.2 we can see that the average number of followers (left) and followees (right) is typically e^{15} and $e^{6.5}$ respectively. Therefore, extracting followers and followees for a single influencer will result in a large network. Even when extracting 0.005% followers ($\approx$ 50000) for a single influencer with parallelisation takes approximately 69 hours.

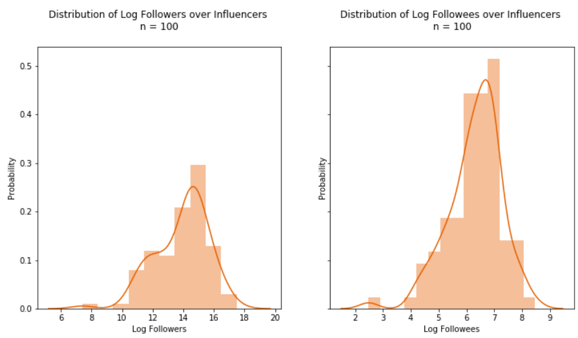

Figure 4.2: Distribution of followers (left) and followees (right) among the 100 identified influencers involved in the analysis. This explains that the size of the network is simply not feasible to extract.

In addition, the explicit following of users is often passive. For example, a user can follow another simply because of a particular post and potentially not interact any further. Therefore, simply following an influencer does not infer idea adoption, and thus authority.

4.1.1 Implicit Connectivity in a User Network

Instead of explicit connectivity (i.e. who follows whom), we will define a new concept of implicit connectivity between users. We define implicit connectivity as interactions between users, which are independent of who follows whom. Therefore, we must implement implicit connectivity at the level of the hashtag conversation. This type of connectivity accounts for active readers who consistently interact with the author, regardless of whether they follow the author or not. We will consider these interactions as readers "lighting up" in response to the author's post at time t.

Although, this way of viewing connectivity between users is a proxy for explicit connectivity, it does appropriately model local connectivity between users and enables us to predict the scale (total number of readers who transmit the information) and strength (the proportion of readers, who consistently light up) of information diffusion in the network. For example, if no users have adopted the hashtag locally within a certain timeframe, then we can expect the scale and strength of information diffusion to be minimal. Furthermore, the granularity of implicit connectivity is finer than for explicit connectivity since we are assessing the degree of interaction at the conversation level. It aims to provide a picture of predictable posting patterns, whereby readers consistently post after an author within the timeframe. This will arguably result in information diffusion to a wider audience.

Furthermore, connectivity between users in social networks can be uncertain and this uncertainty can be caused by external factors, as well as the explicit network itself Han and Li, 2018. They give the example that, at one moment, connectivity between work colleagues can later change when some of those change jobs. Thus, we look to implicit connectivity as a way of expressing this constantly dynamic connectivity between users, for which many models fail to account (Han and Li, 2018).

4.1.2 Group Networks

Since there is no literature which models implicit connectivity between users, we need to derive a new way to model implicit connectivity by measuring the interaction of users in group networks over a hashtag conversation.

Within a hashtag conversation #h, an author posts in the hashtag conversation at time point t. Within a timeframe $t \geqslant T \geqslant t+n$ a set of readers R will post with #h (i.e. "light up"). We will define this implicit connectivity between readers and a specific author within a hashtag conversation as a group network at time t ($g_t^{a,\#h}$) (See Figure 4.3). The scale and strength of implicit connectivity between the author and readers R will be derived from the number times that readers R are involved in group networks with a specific author during the hashtag conversation, where R is propagating the hashtag that the author has used.

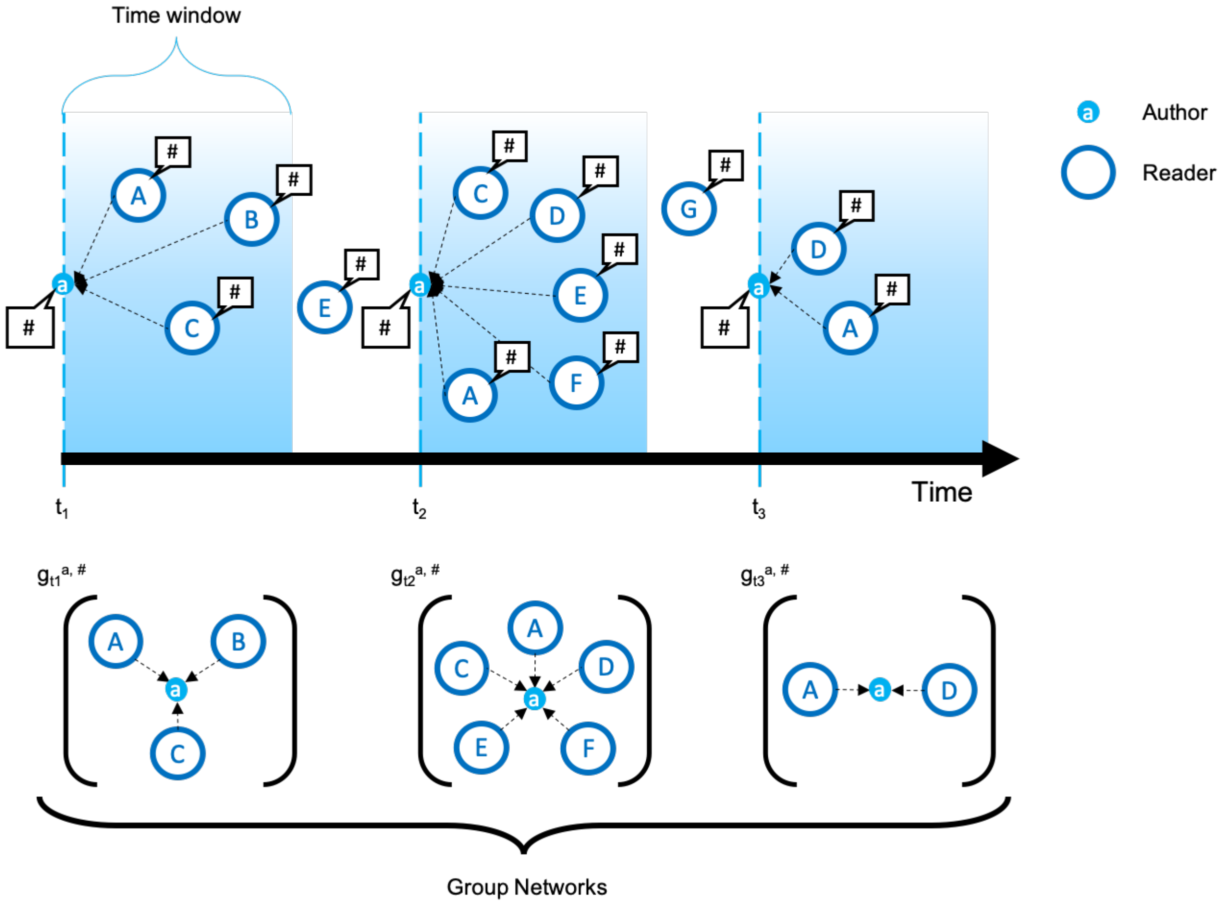

Figure 4.3: Definition of a Group Network for an author at times t_1, t_2, t_3 over a hashtag conversation ($g_{t_1}^{a,\#}$, $g_{t_2}^{a,\#}$, $g_{t_2}^{a,\#}$ respectively)

4.1.3 Adoption Rate

From these group networks, we want to assess the probability that a reader r posts in the hashtag conversation $\#h$ within the timeframe $t \geqslant T \geqslant t+n$, given that the author a has posted in the hashtag conversation : $P(hrR \in t \geqslant T \geqslant t+n|h_a \in t)$ (or $P(h_r|h_a)$ for short). We will term this probability as the adoption rate. We devise the adoption rate by the following:

Over the group networks at time points where author a has posted over the whole hashtag conversation ($g_t \in G$), we count the number of times where a reader appears over all group networks across the whole hashtag conversation. We define this degree of user overlap over group networks as the raw adoption count (See Figure 4.4).

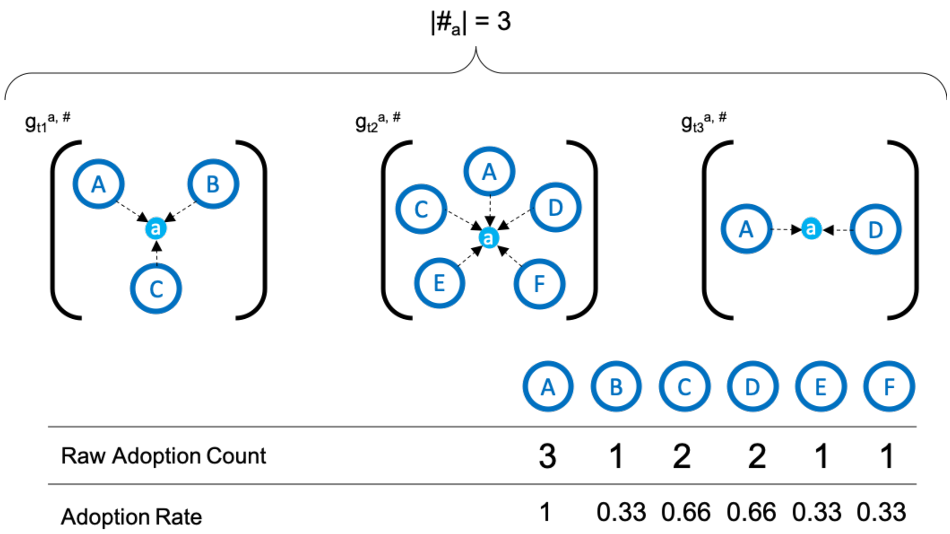

Figure 4.4: Depiction of how we compute the adoption rate, which serves as the basis for assessing authority for a user a over a hashtag conversation #. We firstly count the number of times a reader (A to F in our example) is involved in the group networks of a particular author throughout the hashtag conversation. This is denoted as the raw adoption count. We then normalise the raw adoption counts by the number of posts author a makes in hashtag conversation # to obtain the adoption rate.

This intuition is motivated by the fact that common, overlapping neighbours in a group network equate to direct relationships between users Han and Li, 2018.

This raw adoption count effectively tells us how many times a reader adopts the hashtag after the author.

We then normalise the raw adoption count by the number of posts that user a has authored in the same hashtag conversation to obtain the adoption rate. Thus, we define the adoption rate as follows:

$$P(h_r|h_a) = \frac{|r \in G|}{|h_a|}$$

, where $|h_a|$ is the number of posts author a has posted with hashtag h and $|r \in G|$ represents the number of times where user r is in the group networks over the hashtag conversation (See Figure 4.4).

The adoption rate ranges between 0 and 1, which correspond to no/minimal idea adoption and consistent idea adoption from the author over a hashtag conversation. These adoption rates will be extremely small, and therefore we will apply a log transformation to obtain log adoption rates.

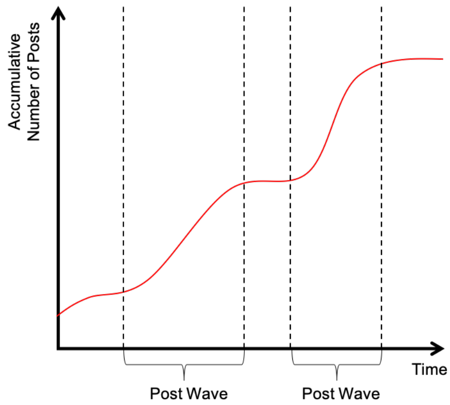

Figure 4.5: Post wave definition and duration

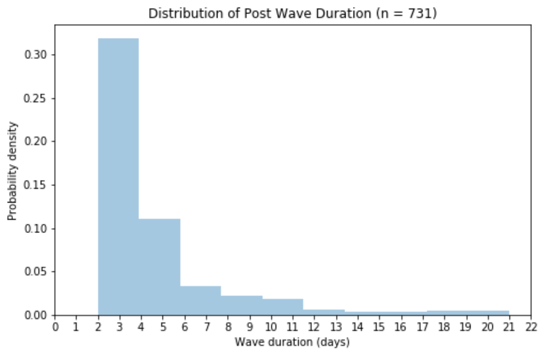

Figure 4.6: Distribution of post wave duration in days over hashtag conversations

To obtain an appropriate time window duration n, we reviewed the distribution of post wave duration in days over hashtag conversations. A post wave (See Figure 4.5) is when we observe noticeable change in the number of accumulative posts in a hashtag conversation, which levels off after a certain period of time. We use a post wave to symbolise a period of heightened posting activity within a hashtag conversation. In our analysis, we measure post wave duration by measuring the distance in days between $\frac{\Delta^2 P}{\Delta T} > 0$ and $\frac{\Delta^2 P}{\Delta T} < 0$, where P represents the accumulative number of posts in a hashtag conversation (See Figure 4.5).

To determine the time window duration to use, we took the top 95th percentile of the distribution of post wave duration (See Figure 4.6), which is 4 days. In our post wave duration, we also want to allow for users posting between post waves, therefore, we multiplied by two, ie. $n = 8$ days.

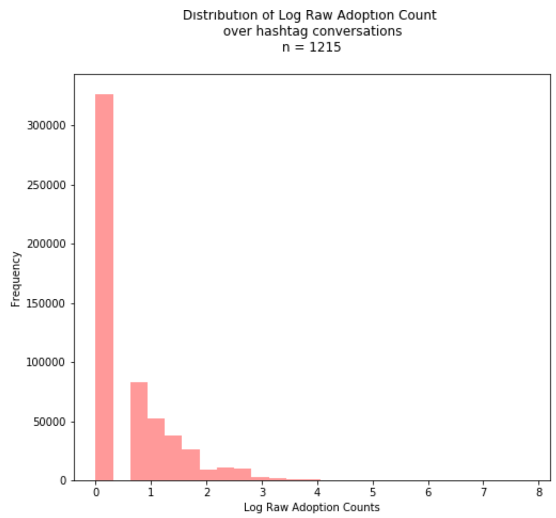

Figure 4.7: Distribution of log raw adoption counts over hashtag conversations.

4.1.4 Internal vs External Exposure

As mentioned earlier, authority is a group process and so can only be observed among those readers in the immediate environment, who have been internally-exposed to the hashtag (Myers et al., 2012, Platow et al., 2015, Watts and Dodds, 2007).

After obtaining raw adoption counts for each author-reader interaction over all hashtag conversations, we need to distinguish between which users are implicitly-connected to the author (internal exposure), against those who are not (external exposure).

In our analysis, we will use the intuition from Myers et al., 2012 to define users being internally-exposed as those who are implicitly-connected to the author. For example, in Figure 4.3, we can say that user A is implicitly connected to author *a*, and hence internally-exposed. We will also define external-exposure as instances where users are not implicitly-connected to the author. For example, in Figure 4.3, we can say that user G is not implicitly connected to the author in any of the group networks, and hence is deemed as externally-exposed.

To separate internally- from externally-exposed users within a hashtag conversation, we use the distribution of the raw adoption counts over hashtag conversations to develop a threshold.

As we mentioned before, raw adoption counts tells us how many times a particular user adopts a hashtag, given that the author has posted that same hashtag.

As we can see from Figure 4.7, the majority of readers appear in no more than one group network with an author over a hashtag conversation. This means that many readers show minimal response to authors overall across hashtag conversations and are thus less likely to have adopted the hashtag from the author and therefore less likely to be implicitly connected to the author.

This intuition serves as the basis for determining the threshold which separates users who are implicitly connected to an author (internal exposure) and the majority of those, who are not (external exposure). We define our threshold as the top 95th percentile over raw adoption counts across all authors and hashtag conversations. This results in a threshold of a raw adoption count = 1 (See Figure 4.7).

For example, in Figure 4.4, we would deem readers A, C and D to be internally-exposed from author a as their raw adoption count is above 1. Readers B, E and F are externally-exposed as their raw adoption count is 1.

We want to ensure that this threshold is appropriate for this analysis. Therefore, we have compared the amount of exposure internally- and externally-exposed readers receive from an author in a hashtag. Here, for the threshold to be appropriate, we expect to see that internally-exposed readers receive greater exposure from authors than externally-exposed readers.

We computed the exposure of an author over a hashtag conversation $E_h^{(a)}$ by the following:

$$E_A^{(h)} = \frac{|h_A|}{|h|}$$

, where $|h|$ refers to the conversation size.

Figure 4.8, shows that internally-exposed readers receive noticeably greater exposure to authors than externally-exposed users ($t = -366.61, df = 563420, p < 2.2e^{-16}$). This suggests that our threshold is appropriate for this analysis.

4.2 Stylistic Accommodation and Mean Attitude Change

From Section 2.2.2, we understand persuasion to be a reader's change in attitude in response to the author's change in attitude of another person. We then went on to argue that attitudes can be encapsulated into linguistic styles Shepard, 2001, which have been represented in either syntactic (Danescu-Niculescu-Mizil et al., 2011, Jones et al., 2014) or semantic space (Nasir et al., 2019).

We also established the importance of discriminating between user homophily and

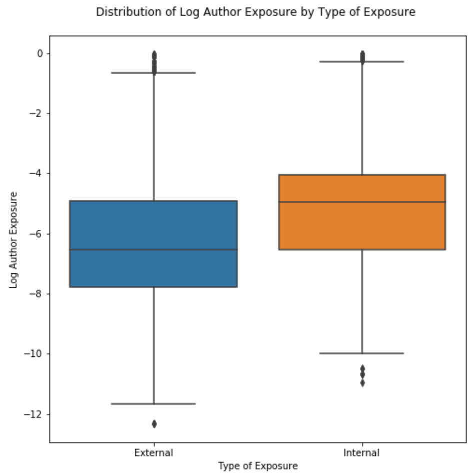

Figure 4.8: Distribution of Log Author Exposure by type of exposure

asymmetric accommodation, whereby homophily is shown by symmetrical changes in attitudes from both the reader R and the author A. Accommodation, on the other hand, is demonstrated by asymmetric changes in attitudes Šćepanović et al., 2017.

This will serve as the principles which have guided the development of the User and Hashtag Stylistic Accommodation functions:

- User Stylistic Accommodation, which measures accommodation relative to users (author and reader) within given timeframe $t-n \geqslant T \geqslant t+n$.
- Hashtag Stylistic Accommodation, which measures accommodation relative to a group / hashtag norm within a given timeframe $t-n \geqslant T \geqslant t+n$.

The aim of these Stylistic Accommodation functions is to discover to what extent influencers (and regular users) are persuasive.

In Section 2.2.2, we also looked at models that explain the routes by which influencers can be persuasive. We described the central and peripheral routes to persuasion from the Elaboration Likelihood Model (ELM) Petty and Cacioppo, 2012. The central route, which utilises logic and reasoning to persuade readers is argued to show minimal changes in attitude over time among readers after a persuasion attempt (i.e. an influencer's post). The peripheral route, which relies on superficial features (e.g. influencer attractiveness) to persuade readers proposes wide variances in attitudes.

To quantify the variation in attitude over a hashtag conversation, we will introduce Mean Attitude Change, with the aim of assessing which route to persuasion influencer's utilise over readers. This is detailed in section 4.2.4.

4.2.1 Stylistic Accommodation

Both User and Hashtag Stylistic Accommodation functions build on the principles of the Zelig quotient (Jones et al., 2014), the Word Mover's Distance (WMD) formulation (Nasir et al., 2019) and the CAT model theory (Shepard, 2001) as explained below.

Before we start defining the functions, let us introduce the key terminology, which follows that used in Jones et al., 2014.

a The author's post: This is the Instagram post which is published by the author at time t.

r The reader's response post: This is the post which responds to post a within a certain time window $t-n \geqslant T \geqslant t+n$.

μ_R The reader's baseline style: The baseline style is defined as the typical linguistic style that the reader R assumes when posting content. The baseline style for a user U is computed as follows:

$$\mu_U = \frac{embed(P(U))}{|P(U)|}$$

, where $embed(p)$ represents the doc2vec representation of post p in semantic space and $|P(U)|$ refers to the total number of posts user u has authored.

μ_H The hashtag norm: The hashtag norm represents a group norm, a shared linguistic style across a hashtag conversation. The hashtag norm is computed as follows:

$$\mu_H = \frac{embed(P(U \in H))}{|P(U \in H)|}$$

, where $|P(U \in H)|$ represents the total number of posts in the hashtag conversation (conversation size).

All of the linguistic style embedding vectors mentioned above are normalised, such that the length of the vectors $||\vec{v}|| = 1$. We do this because we are focused on the direction of these vectors, rather than the magnitude of the vectors used to compute stylistic accommodation.

While the principle of the Zelig quotient allows to assess the degree of accommodation, its main issues for our purposes are as follows:

1. Unlike the WMD formulation, the quotient does not represent linguistic style in semantic space, which was argued to better represent changes in attitudes than syntactic space (See Section 2.2.2).

2. Along with all formulations discussed in Section 2.2.2, the Zelig quotient does not differentiate accommodation from homophily.

We need define a new approach to measure change in linguistic style over our hashtag conversations, using the some of the principles from the formulations stated earlier.

To obtain linguistic styles in semantic space, we trained a Doc2Vec model with 70,081 randomly chosen posts in English, across all extracted hashtag conversations.

A Doc2Vec model is an unsupervised technique, based on a neural network, that is used to translate documents (in our case, individual posts) into a n-dimension vector in semantic space.
As with a Word2Vec model, in a Doc2Vec model (as per Le and Mikolov, 2014), the goal is to output the context in which the input word appears (See Figure 4.9). Here, the context refers to a specific number of tokens, which appear after the input word. During the training of the model, the words are represented as one-hot encoded vectors with dimensionality $d = |V|$ where $|V|$ is the vocabulary size.

The architecture of the model, is made up of an input, hidden and output layer, where the hidden layer of n dimensionality represents a word embedding of the input word. In our analysis, we refrained from choosing a large n, firstly due to a relatively small vocab size in the posts ($|V| = 406,254$) and to avoid sparsity. Therefore, we chose $n = 64$ as a reasonable compromise. From the individual word embeddings, we sum them up over the document to form a document embedding. Using the example in Figure 4.9 the one-hot encoding for "flashback" is used as input with the target output being "to a lovely" (the next 3 tokens in the sentence). The input and target output is used to train a 64-dimension hidden layer. This is summed for all words in the post to define a document embedding.

To identify English posts, we utilised a Multinomial Naïve Bayes natural language classification model, which reports high accuracy over classifying texts into one of 21 European languages ($> 99\%$) (Vasilynonok, 2018). From Figure 4.10, we see that overall over 80% posts are in English, which motivates the use of only English posts in the training set.

Before training, we also cleaned the individual posts by removing any user mentions (eg. @michelle_lewin), punctuation, and English stopwords, which do not contribute to the overall meaning of the post. To determine the context window size, we assess the distribution of cleaned post length, whereby the top 2nd percentile is chosen to ensure that the context size is small enough to obtain contexts from 98% English

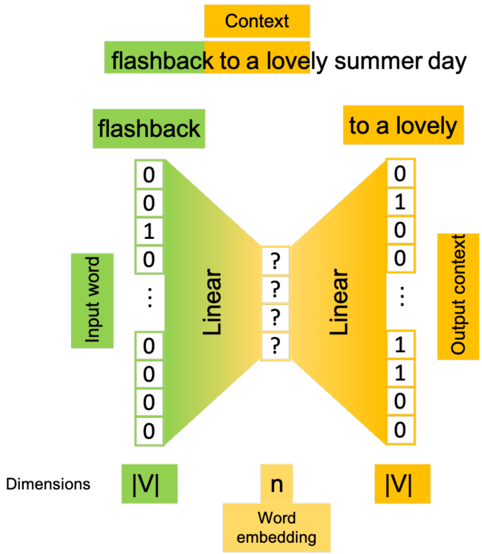

Figure 4.9: The architecture of Doc2Vec model and how the Doc2Vec model is trained over the example sentence.

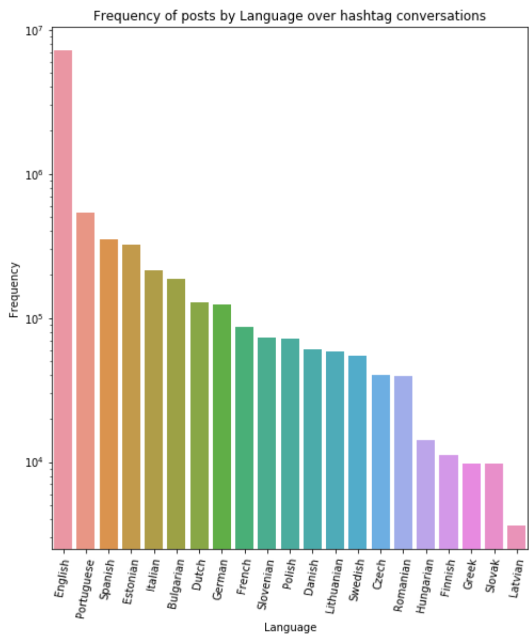

Figure 4.10: Distribution of posts by language, which have been identified from the natural language classification model (Vasilynonok, 2018)

posts. The resulting context window size is 3 tokens.

Taking inspiration from CAT theory's decomposition of accommodation, we will break down the stylistic accommodation formulations into:

- The extent of accommodation, which computes accommodation degree (as out-line above)

- Symmetry, which includes accommodation direction. The symmetry is also responsible for indicating the degree of homophily between the author A and reader R.

We also need to consider the symmetry of the linguistic style change to establish who is accommodating/moving to whom. Again, there is no formulation which we can apply from the current literature, so we have built a suitable stylistic accommodation formula, based on relevant academic literature outlined in Section 2.2.2.

From CAT, the premise of accommodation overall rests on the symmetry/asymmetry of the movement from a readers's baseline $\vec{\mu_R}$ to their recent post $\vec{r}$ ($\vec{\mu_R r}$) in comparison with the author (See Figure 4.11). We will refer to ($\vec{\mu_R r}$) as the change is user R's linguistic style within a timeframe $t - n \geqslant T \geqslant t + n$.

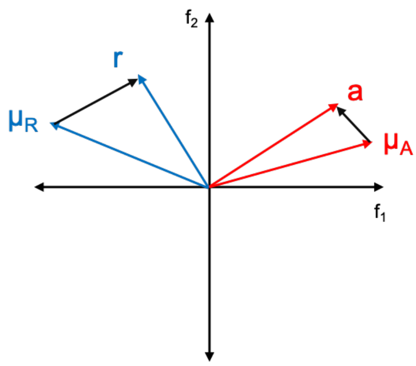

Figure 4.11: The premise behind accommodation is measuring the symmetry/asymmetry of the movement from a reader's baseline $\vec{\mu_R}$ to their recent post $\vec{r}$ ($\vec{\mu_R r}$) in comparison with the author

To obtain the time window n, for all users, we obtain posting intervals (the time gap between a user's posts) for each user. From the data, the top 95% of users post within a 28-day period. From the accommodation premise, we need to compute user R's baseline style before user A's post and get R's most recent post after user A's post. To increase the chances of obtaining the same user R, both before and after a particular user's post, we use ($n = 28$ days) as the time window to compute stylistic accommodation.

4.2.2 User Stylistic Accommodation

As mentioned earlier, user stylistic accommodation is used to measure the degree of a reader's change in linguistic style, given the author's change in linguistic style at time t. This form of accommodation looks at the author's degree of persuasion between the author and the reader, irrespective of the group / hashtag topic norm.

The extent of accommodation accounts for the amount of change in a user's linguistic styles (degree). We measure reader's extent of accommodation as the following:

$$T_{R \longrightarrow A} = \frac{||r\vec{\mu}_A||}{||r\vec{\mu}_A \cdot \mu_R\vec{\mu}_A||}$$

The vector $||r\vec{\mu}_A||$ describes how far the reader's post r is from the author's baseline style μ_A.

For the denominator, from $a \cdot b = ||a|| \cdot ||b||cos(\Theta)$, we can say that the dot product between $r\vec{\mu}_A$ and $\mu_R\vec{\mu}_A$ essentially multiplies the length of these together (along with the cosine of the angle between them). Therefore, the extent function essentially finds how much $r\vec{\mu}_A$ in relation to $r\vec{\mu}_A$ and $\mu_R\vec{\mu}_A$. Therefore, extent ranges from 0 to $+\infty$. $T_{R \longrightarrow A} = 0$ indicates that the reader remains at their baseline style and therefore has not moved towards the author's baseline. $T_{R \longrightarrow A} = +\infty$ means that either: the reader's post is the exact same as the author's baseline style or the author's and reader's baseline linguistic styles are the exact same. Both of these indicate significantly extreme extents of accommodation.

For symmetry, we have to compare the changes in both the author and reader's linguistic style, as well understanding whether they are converging together or not. Therefore, we split the symmetry computation accordingly. Symmetry between author A and reader R is measured by the following:

$$S_{RA} = (||\vec{\mu_R r}|| - ||\vec{\mu_A a}||) \times (1 + D(sim_A, sim_R))$$

, such that

$$sim_A = D(\vec{\mu}_A \cdot \vec{a}) \quad \text{and} \quad sim_R = D(\vec{\mu}_R \cdot \vec{r})$$

$$D(\vec{x}, \vec{y}) = \frac{(x \cdot y) + 1}{2}$$

The first part of the equation $(||\vec{\mu_R r}|| - ||\vec{\mu_A a}||)$, which compares the author's and reader's movement, is simply done by subtracting the euclidean distances of their respective movements from each other. Therefore, we expect symmetrical movement to produce 0. If this part of the equation is greater than 0, then we can say that the reader is asymmetrically accommodating. when this part of the equation < 0, then the author is asymmetrically accommodating.

The second part of the equation $(1 + D(sim_A, sim_R))$ assesses whether both the reader and author's movement converges towards each other or not. As stated in 2.2.2, symmetrical convergence of linguistic styles indicates homophily.

The dot function D is in place to ensure that the function returns values between 1 and 0, where 1 signifies that x and y are of the same direction and 0 means that x and y are in pointing opposite directions.

The sim_A and sim_R firstly quantifies the author's and reader's movement in attitudes relative to their respective baseline styles. Here, they will return 1 if they are close their baseline linguistic style and 0 if they are dissimilar. Therefore, if these movements are symmetrically converging then, we expect $D(sim_A, sim_R) \to 1$ and 0 to indicate that the author and reader are not converging.

Thus, $S_{RA} = 0$ indicates symmetry in the author and reader's changes in attitudes and that they are converging. $S_{RA} > 0$ indicates asymmetric change in attitudes, such that $R \to A$ more than $A \to R$. $S_{RA} < 0$ indicates asymmetric change in attitudes in the other direction.

From the extent and symmetry, we compute User Stylistic Accommodation between the reader R and author A as follows:

$$Acc_{RA} = \sigma(S_{RA} + \frac{T_{R \to A} - T_{A \to R}}{||T_{R \to A}|| + ||T_{A \to R}||}$$

, where

$$\sigma(x) = \frac{2}{1 + e^{-x}} - 1$$

We include a sigmoid function to collapse the range to -1 and 1, which correspond to extreme $R \to A$ accommodation and extreme $A \to R$ accommodation respectively. The function also produces 0, which corresponds to homophily (symmetrical movement of author and reader towards each other).

The extents (for author and reader) range from 0 to $+\infty$. Passing these large extent values through the above function would skew the final accommodation result to the extreme values (1 or -1) so we need to normalise the extents.

4.2.3 Hashtag Stylistic Accommodation

The hashtag stylistic accommodation also measures the degree of a change in a reader's linguistic style, given the author's change in linguistic style at time t. However, unlike user stylistic accommodation, hashtag stylistic accommodation is dependent on a group / hashtag norm.

From Hogg and Reid, 2006 we can define a hashtag norm as a regularity in behaviour over the social group. Therefore, we will define a hashtag norm to be the average linguistic style over all posts in the hashtag conversation. From CAT, we will assume in our analysis that users involved in a given hashtag conversation form a group, which define a norm (Shepard, 2001).

To compute the extent of accommodation relative to the hashtag norm, we now need to consider the amount of change in a user's linguistic styles (degree) with respect to how close the user's baseline is to the hashtag norm. Therefore, we measure readers extent of accommodation, relative to the hashtag norm μ_H as the following:

$$T_{R \to H} = \frac{||r\vec{\mu}_H||}{||r\vec{\mu}_H \cdot \mu_R\vec{\mu}_H||}$$

Here, the motivation for this formulation is motivated in the User Stylistic Accommodation, which now replaces the author A with the hashtag conversation H.

As before, extent ranges from 0 to $+\infty$. $T_{R \to H} = 0$ indicates that the reader remains at their baseline style and and therefore has not moved towards the hashtag norm. $T_{R \to H} = +\infty$ means that either: the reader's post is the exact same as the hashtag norm or the reader's baseline linguistic style is the exact same as the hashtag norm. Both of these indicate significantly extreme extents of accommodation.

For symmetry, we compare relative changes in user's linguistic style, as well as the direction of accommodation, relative to the hashtag norm. We compute hashtag symmetry between author A and reader R by the following:

$$S_{RA}^{(H)} = (D(\mu_R, \mu_H) - D(\mu_A, \mu_H))$$

We interpret the symmetry of accommodation in the same manner as in User Stylistic Accommodation.

From the extent and symmetry, we compute Hashtag Stylistic Accommodation between the reader R and author A as follows:

$$Acc_{RA} = \sigma(S_{RA} + \frac{T_{R \to H} - T_{A \to H}}{||T_{R \to H}|| + ||T_{A \to H}||}$$

, where

$$\sigma(x) = \frac{2}{1+e^{-x}} - 1$$

As motivated in User Stylistic Accommodation, we use the sigmoid, so that Acc_{RA} ranges between -1 and 1, which correspond to extreme $R \rightarrow A$ accommodation and extreme $A \rightarrow R$ accommodation relative to the hashtag norm respectively. The function also produces 0, which corresponds to homophily (symmetrical movement of author and reader towards the hashtag norm). Extent normalisation is applied as for User Stylistic Accommodation.

4.2.4 Mean Attitude Change

In this subsection, we will introduce Mean Attitude Change, which evaluates the reader's variation of attitude over a hashtag conversation H, after the influencer's persuasion attempt. We define an influencer's persuasion attempt as a post which occurs after the reader's first post in the hashtag conversation.

From the ELM Petty and Cacioppo, 2012, the key feature which discriminates between the central and peripheral routes is the variation in attitude over time after a persuasion attempt. The central route is denoted by low attitude variation and peripheral high attitude variation.

Therefore, the goal of Mean Attitude Change is to measure attitude variation by measuring variation in a reader's linguistic style over a hashtag conversation H, provided that they have been "persuaded" by an influencer in H. We do this to ensure that we assess attitude change after interaction with influencers.

We measure variation relative to the reader's first post in the hashtag after the influencer's persuasion attempt to understand how the attitude changes after they are "persuaded" by the influencer. Therefore, we can assess whether readers were persuaded via the central or peripheral route.

From this, we can take the standard deviation of a reader's linguistic style over a hashtag conversation H with respect to the first post the reader published in the hashtag conversation after the influencer's persuasion attempt ($P(R)_1 \in H$), which is computed as follows:

$$\sigma_R^H = \sqrt{\frac{1}{|P(R) \in H|_{T>t}} \sum_{i \geqslant t}^{|P(R) \in H|} (embed(i) - embed(P(R)_1 \in H))^2}$$

, where $|P(R) \in H|_{T>t}$ represents the total number of posts the reader has published in the hashtag conversation H after the influencer's persuasion attempt at time t. (See Figure 4.12).

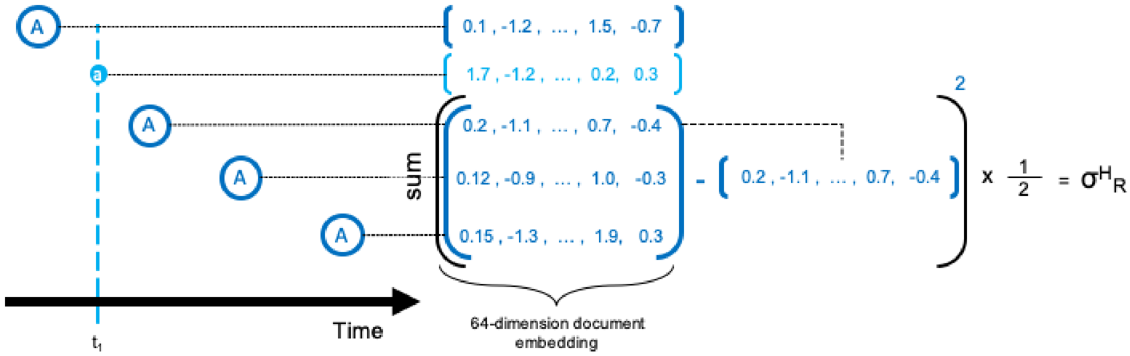

Figure 4.12: Visual description of how we compute σ_R^H for a reader R in a hashtag conversation H ("A" denotes the reader before and after the author posted)

This will show variation over the 64 semantic dimensions with respect to the first post that the reader made in the hashtag conversation after influencer's initial persuasion attempt. Currently, this is hard to interpret and compare against other readers.

To facilitate comparison between readers, we collapse the resulting 64-dimension σ_R^H into a scalar value by computing the mean over the dimensions to gain an insight to understand the average variation in linguistic style, hence Mean Attitude Change:

$$\lambda_R^H = \frac{1}{D}\sum_{d=1}^{D} \sigma_{R\,d}^H$$

Here, we deem $\lambda_R^H = 0$ as no variation in the reader R's linguistic style over the hashtag conversation H, and thus indicates the reader's attitude does not deviate from their initial attitude in the hashtag conversation, and hence indicates that the influencer used a central route to persuasion. $\lambda_R^H > 0$ indicates greater variation in the reader's linguistic style, which relates to a more peripheral route to persuasion.

To obtain a range of λ_R^H, we need to compute the maximum λ_R^H. To compute the maximum λ_R^H, we identified the maximum and minimum post embedding values for each dimension, which in general are in the range of -2.77 and 2.98. We then computed the Mean Attitude Change over these two vectors to understand the maximum Mean Attitude Change. The maximum $\lambda_R^H = 1.49$, which indicates a peripheral route to persuasion.

4.3 Topic Dominance

This section describes the topic dominance function, based upon the sparse literature outlined in Section 2.2.3. From this, we independently developed this function to capture how influencers participate or respond to trends (i.e. whether they set trends or 'ride the wave' and simply follow them), which is associated with the topic initiation and uptake dimension of influence.

For this function, we will follow the intuition from Shen et al., 2017, where we will explore how influencers respond to trends over time. Therefore, we introduce topic dominance as a way to describe the likelihood that an author was involved in a trend (which we will observe over hashtag conversations) or initiated a new trend, evidenced by uptake at a certain speed and magnitude.

Before we get into the details of the function itself, it is worth noting the following:

- We are computing Topic Dominance by looking at the increases in the accumulative post count of the hashtag topic H with respect to time T ($\frac{\Delta H}{\Delta T}$). We do this so that we can robustly trace uptake as an increase in conversation size over a certain time window.

- We are assessing Topic Dominance for an author A, who posts at a given time t. To understand whether the author is following a trend (i.e. "riding the wave") or initiating a new trend, we need to observe this uptake over a time window $t-n \geqslant T \geqslant t+n$. We used the same time window identified in Section 4.1 of $n = 4$ days (See Figure 4.6) to accommodate for a wave taking place in the time window.

As a basic starting point, we can compute uptake from a user's post at time t simply as:

$$\frac{H_{t+n} - H_t}{T_{t+n} - T_t} \quad \text{or} \quad \frac{\Delta H_n}{\Delta T_n}$$

We will refer to this as P_2'. Although this allows us to see the change in number of posts after the author has posted, it does not provide robust evidence of whether the author is riding the wave or initiating a trend because it does not show whether the uptake was driven by the author or by another post published in $t-n \geqslant T \geqslant t$.

A solution to this issue is to compare P_2' against $\frac{H_{t+n}-H_{t-n}}{T_{t+n}-T_{t-n}}$ or $\frac{\Delta H}{\Delta T}$ (we will refer to this as P_1'). From this we can ascertain whether the uptake directly follows the author's post at time t or the author is simply riding the wave.

So far we can now differentiate whether the author is riding the wave or initiating a new trend within a hashtag conversation within a time window. However, we so far do not have an idea of the rate of uptake after the author's posts $t \geqslant T \geqslant t+n$. Therefore, we need to incorporate $\frac{\Delta^2 H_n}{\Delta^2 T_n}$, which we will refer as P''. If P'' is high, then this suggests that overall change in H occurs after a post published after the author's. A low P'' suggests that the increase in H directly follows the author's post. Since time T in our analysis is discrete, we compute P'' by the following:

$$P'' = \frac{\Delta^2 P}{\Delta^2 T} = \frac{P_{t+2} + P_{t-2} - 2P_t}{4h^2}$$

, where P_t represents the accumulative number of posts at time t and h represents a regular time interval (method taken from "Numerical (Second) Derivative of Time Series Data", 2018).

P'' only denotes the uptake trajectory, but does not give firm evidence to suggest that the uptake directly follows the author's post at time t. We can remedy this by including the proximal gradient (P_{prox}). For uptake to directly follow after the author's post, we would then expect P_{prox} or $\frac{\Delta H}{\Delta T} \to +\infty$ where $\Delta T \to 0$.

Figure 4.13 shows all the necessary gradients required to compute Topic Dominance.

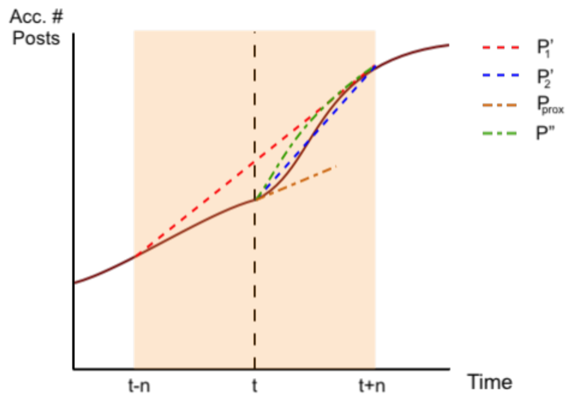

Figure 4.13: The components involved in computing topic dominance

We combine these uptake gradients to compute an author's Topic Dominance at a given time point t over a hashtag conversation as the following:

$$TD_t^A = \sigma(\Delta)$$

, such that

$$\Delta = \frac{P'_2}{P'_1} \times P'' \times (1 + P_{prox})$$

and

$$\sigma(x) = \frac{1}{1 + e^{-x}}$$

The TD function ranges between 0 and 1. $TD_t^I = 0$ means that there is no uptake after the influencer's post at time t, which means that the author has not started a trend. $TD_t^I = 0.5$ means that the author is simply "riding the wave" and following a previously-initiated trend at time t. In Figure 4.14, we see varying instances within a post wave where $TD_t^I = 0.5$. For example, $TD_2^I = 0.5$ because the rate of uptake after the influencer's post $P'' = 0$, and therefore $\Delta = 0$ and thus $TD_2^I = \sigma(\Delta) = \sigma(0) = 0.5$. TD_3^I, TD_4^I, follow the same logic. In the case of $TD_1^I > 0.5$, we see that the influencer has posted at the beginning of a post wave, which indicates that they has initiated a trend.

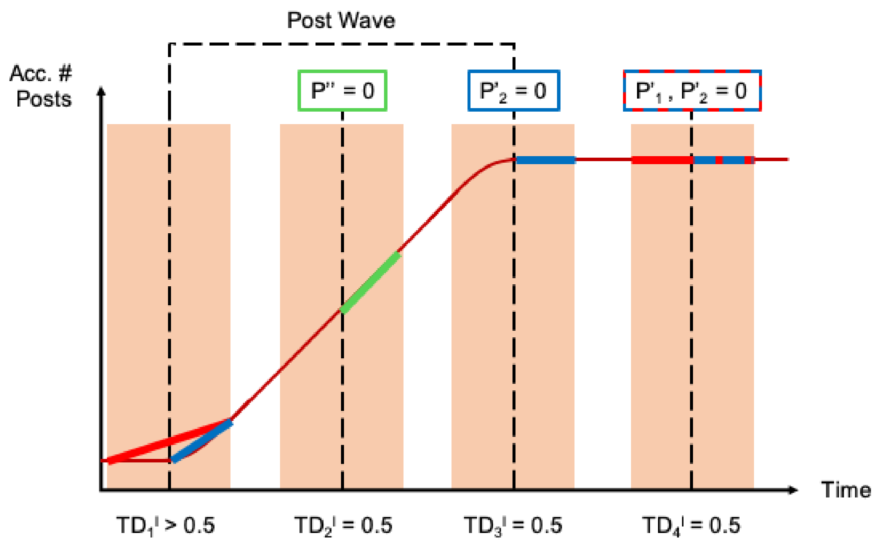

Figure 4.14: Visualisation of Topic dominance at various points of a post wave with explanations as to why some points in the post wave will result in $TD_t^I = 0.5$

4.4 Hypotheses

In this section we will formalise the hypotheses to be analysed. Firstly, we will establish whether Instagram influencers are significantly different from regular users in their behaviours. Then we will explore patterns of influential behaviour using the dimensions of Influence Space described in section and the derived formula and measures defined earlier in Sections 2.1 and 4 respectively.

The hypotheses are as follows:

1. Differentiated Authority: We expect influencers to be more authoritative in hashtag conversations than regular users. We also expect influencers to have larger audience sizes over hashtag conversations in comparison with regular users, with a greater proportion of the audience involved in group networks over the hashtag conversation.

   This can be split into the following sub-hypotheses:

   1.1 We expect that there are a greater proportion of readers who are internally-exposed by influencers than regular users. We also expect influencers to have larger audience sizes over hashtag conversations in comparison with regular users, with a greater proportion of the audience involved in group networks over the hashtag conversation.

   1.2 Over internally-exposed readers, we expect the average adoption rate to be higher for influencers than regular users.

   1.3 We expect influencer domain audience sizes and network group sizes to vary widely by domain however network group size (the active user base) is broadly proportional to audience size. This relates to an influencer's need to build an active core group as they extend their audience, to sustain their authority.

   1.4 We expect influencer's authority to be domain- and topic-dependent rather than universal.

   1.5 We expect influencers to have greater authority over readers who exclusively respond to posts from influencers of a single domain rather than across multiple domains.

2. Differentiated Persuasion: We expect influencers to be more persuasive in hashtag conversations than regular users.

This can be split into the following sub-hypotheses:

2.1 We expect stylistic accommodation between influencer-user interactions to be greater on average than stylistic accommodation in user-user interactions.

2.2 We expect influencers to have high situational influence over a small proportion of conversations and low global influence across their domain (Rosenthal and McKeown, 2016). Therefore, we expect persuasion to be domain- and topic-dependent.

3. Persuasion Through Similarity: From Rosenthal and McKeown, 2016, we expect to also see similarity between influencers and regular users yield high observed influence. Cartwright, 1951 suggests that to change the opinion of another individual, one must have a strong sense of belonging to a group. Therefore, the more an individual reflects the norm within the group, then the more likely they are to influence them (Platow et al., 2015, Turner, 1991).

   This can be split into the following sub-hypotheses:

3.1 We expect influencers on average to be closer to hashtag norms than users.

3.2 We expect the influencer to have close proximity to the hashtag norm than regular users.

4. Persuasion Through Style or Substance. We expect influencers to predominately use the peripheral route to persuasion over central route. In the case of Instagram influencers, we predict that, given the visual nature of Instagram, influencers typically persuade readers via the peripheral route than the central route.

5.1 Topic Dominance Through Trend Setting: We expect that influencers initiate more trends than regular users. Therefore, we expect that influencers have a greater proportion of $TD > 0.6$ than users.

5.2 Topic Dominance Through Niche Domain Focus. We expect influencer trend initiation to be domain-dependent.

Chapter 5

Analysis

In our analysis, we will use $\alpha = 0.05$ to determine statistic results as significant.

5.1 Authority

Hypothesis 1.1 - There are a greater proportion of readers who are internally exposed by influencers than regular users

Figure 5.1 shows the proportion of readers who are internally (by direct authority) and externally (from other sources) exposed by user type (influencer and regular users). Here, we can see that overall, for both influencers and regular users, over 50% readers are externally-exposed.

Influencers (I) have approximately 10% more internally-exposed readers than regular users (U).

For the reasons explained in Section 4.1, we exclude externally-exposed readers when evaluating an author's level of authority.

As we can observe from Figure 5.2, there is an overall positive association between average log group size and log audience size for both influencers and users. Influencers do have noticeably larger average group sizes than users at each discrete audience size interval. We cannot conclude however that the variation in log group size can be explained by user type themselves ($F[1,988] = 2.29, p = 0.131$).

Figure 5.3 shows the mean audience proportions by the same log audience size intervals. We compute mean group network/audience proportion by

$$\frac{\mu_{|G_U^H|}}{mu_A^H{}_U}$$

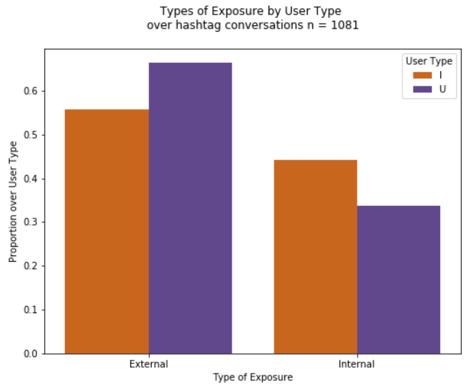

Figure 5.1: Type of Exposure by User Type

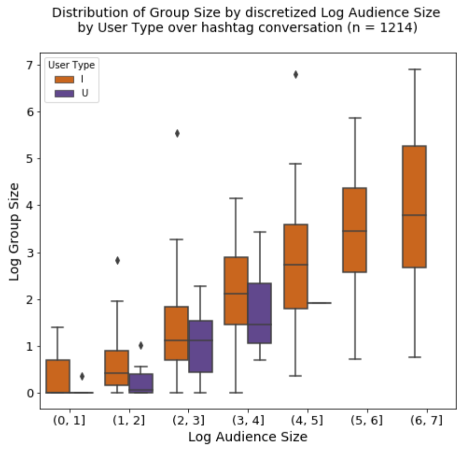

Figure 5.2: Distribution of log group size by discretized log audience size to show the positive association between group size and audience size among influencers

, where A_u^H corresponds to the audience size in hashtag conversation H for user $u \in U$ and $|G_u^H|$ represents the group network sizes over hashtag conversation H for user $u \in U$.

We notice an overall average decay in average group network/audience proportions as the log audience itself increases. We also observe that, for influencers, the rate of average group network/audience proportion (proportion of active interactants in the audience) decay is slower than that among regular users ($\beta = 0.653$ [$R^2 = 0.992$] and $\beta = 0.514$ [$R^2 = 0.954$] respectively).

Overall, the results in this section show influencers have a greater proportion of internally-exposed readers than the regular users. Influencers on average have larger group network sizes over hashtag conversations, in comparison with regular users.

Compared to regular users, influencers also show a slower decay in average group

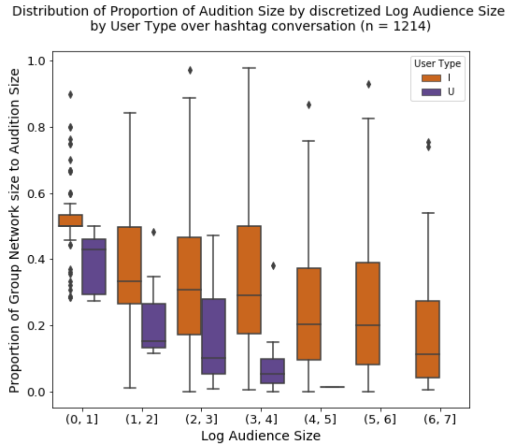

Figure 5.3: Distribution of average audience proportion over discretized audience sizes by user type. The average audience proportion is computed as $\frac{\mu_{|G^H_U|}}{mu^H_A U}$

network/audience proportion as the audience size increases.

Hypothesis 1.2 - Average adoption rate higher for influencers than regular users

Figure 5.4 shows that, overall, influencers have significantly higher average log adoption rate over hashtag conversations, in comparison with regular users ($t = -77.16, df = 188160, p < 2.2e-16$).

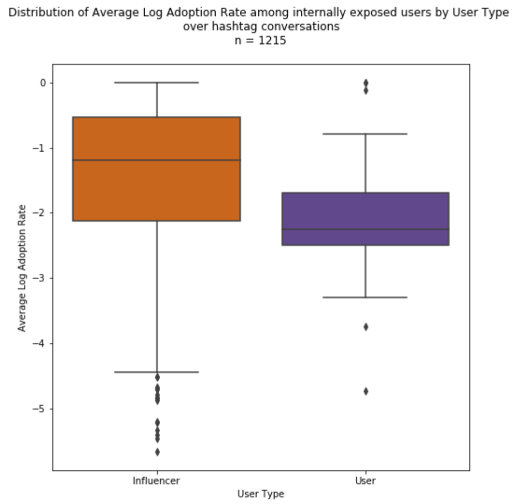

Figure 5.4: Distribution of average log adoption rate among internally exposed users over hashtag conversations by user type

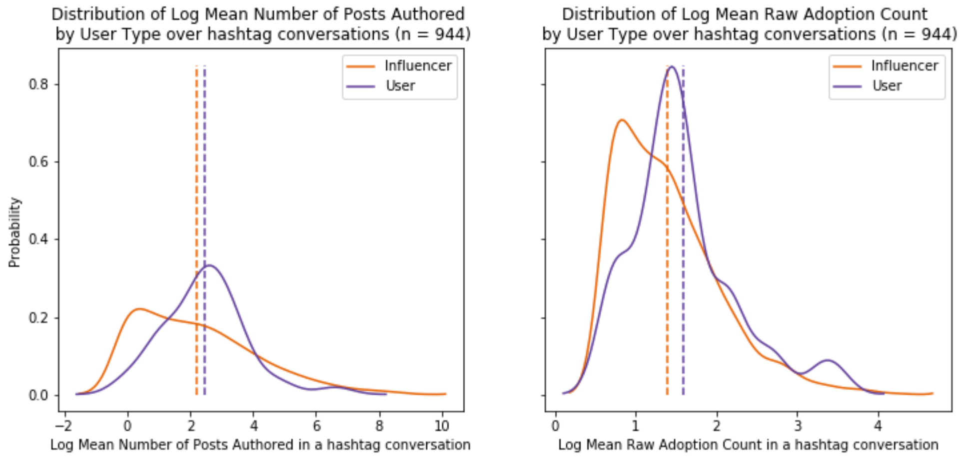

Figure 5.5: (Left) The distribution of mean number of posts authored by user type over hashtag conversations (Right) Distribution of mean raw adoption counts by user type over hashtag conversations. The dashed lines in both plots represent the user type's mean of the respective distributions

However, this result could be biased due to the volume of posts authored by influencers compared to regular users so we need to explore relative Instagram post volumes. In Figure 5.5 (left), we observe that the average influencer's log mean number of posts over hashtag conversations is marginally lower than regular user's.

In Figure 5.5 (right), the distribution of reader raw adoption counts for influencers is marginally lower on average than regular users, however, this difference is not significant ($t = -1.1748, df = 41.483, p = 0.247$).

In summary, influencers show significantly greater log adoption rate than regular users over hashtag conversations. We also observe regular users typically post more than influencers over hashtag conversations and, as a consequence, the readers of regular user posts show a greater log raw adoption counts than influencers.

Hypothesis 1.3 - Influencer's audience sizes and network group sizes vary widely by domain however group size is broadly proportional to audience size

From Figure 5.6, we notice a similar observation, noted in section Hypothesis 1.1, in that, for the majority of domains, more than half of readers appear to be externally-exposed. By definition, these readers are excluded from the rest of our analysis.

Fashion influencers show the highest proportion of internally-exposed users ($>$

Figure 5.6: Type of Exposure by Domain across influencers and hashtag conversations.

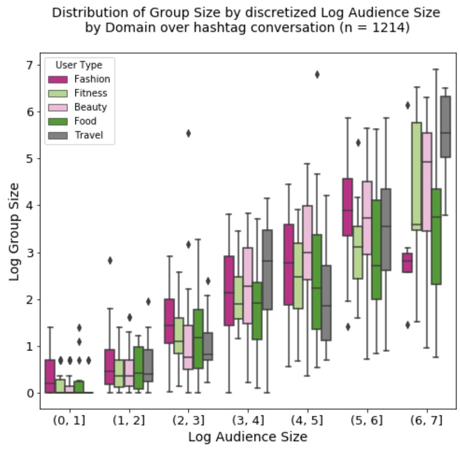

Figure 5.7: Distribution of group network size by domain over log audience sizes

50%) relative to the other four domains. Travel influencers have the lowest proportion of internally-exposed users ($\approx 35\%$).

From Figure 5.7, we see that the log group network size (the number of users who respond to a post from an author a at time t) increases linearly as the audience size increases across all 5 domains (See Table 5.1).

Fashion influencers consistently have larger group network sizes over hashtag conversations, in comparison with the other domains. However, we notice a sudden drop in average group size for log audience sizes between e^6 and e^7.

From Figure 5.8 we observe the growth in log audience size relative to the group network/audience proportions (the strength of authority), for the majority of log audience size intervals. Fashion influencers again see greater proportions of their audience in their group networks (i.e. high degree of active conversations), in comparison with fitness or food, for example.

Domain	Slope	R^2
Beauty	-0.710	0.984
Fashion	-0.571	0.896
Fitness	-0.671	0.980
Food	-0.567	0.991
Travel	-0.805	0.897

Table 5.1: The log intercepts and slopes which show the rate of the average decay of group network size / audience proportion over an increasing audience size. The high R^2 across all domains indicates that the slope and intercepts are relatively accurate.

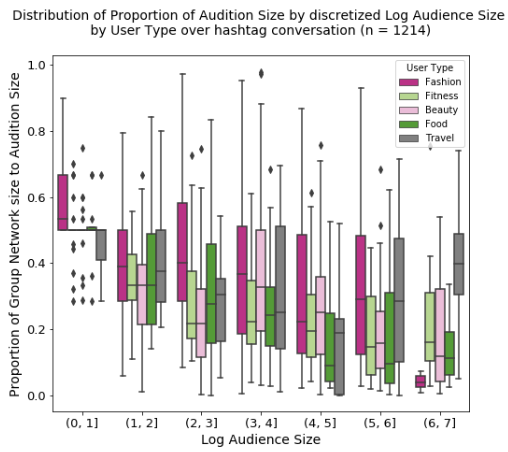

Figure 5.8: Distribution of average audience proportion over discretized audience sizes by user type. The average audience proportion is computed as $\frac{\mu_{|G_U^H|}}{mu_{A\,U}^H}$

For larger audience sizes ($e^5 - e^7$), travel influencers demonstrate greater proportions of their audience are involved in their group networks over hashtag conversations, even more than for fashion influencers.

In summary, there is an overall pattern that as the audience size increases, so does the group network size. There is also is relatively large degree of variance in group network size over domains, particularly as the audience sizes are large ($e^6 - e^7$).

In terms of the group network size / audience proportion, we observe an overall decline in this as log audience size increases across the majority of domains (except travel, in particular, where the proportion noticeably grows from large audiences [$e^6 - e^7$]). We also observed that there is variation in the decay of the group network size / audience proportion as the log audience size increases.

Hypothesis 1.4 - Influencer's authority is domain and topic-dependent

To assess whether domain and/or hashtag topic have a significant effect on the log adoption rate, we need to observe whether the log adoption rate varies across domains and hashtag topics separately. Therefore, we will break down the analysis by domain-level authority, which spans over the influencer domains discussed in Section 3 to check if log adoption rate varies across influencer domains. We will subsequently review domain- and topic-level authority, which will focus on hashtag conversations that occur over all five influencer domains. Only 10 hashtag conversations involved participation from influencers over all 5 domains.

Figure 5.9 shows average log adoption rates by domain. We can see noticeable variability, for example, food influencers demonstrate a higher average log adoption rate from their post readers, on average, than for fashion and fitness influencers.

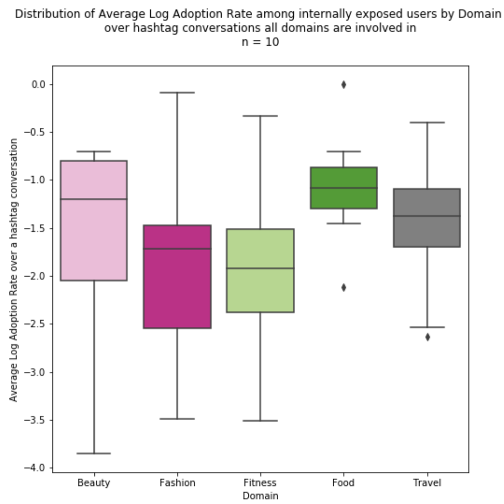

Figure 5.9: Distribution of average log adoption rate by influencer domain over the subset of hashtag conversations with internally-exposed users

From Figure 5.9, we can see that fashion and fitness influencers on average post more in a hashtag conversation and experience a greater mean raw adoption count, on average, in comparison with the other domains. This can potentially mark a negative association between log adoption rates and their parameters: the higher the log adoption rate, the less the influencer posts in a hashtag conversation and the lower the raw adoption count on average.

From Figure 5.11 we also observe similar variability of log adoption rate over hashtag topics. We have selected only the hashtags topics that are common to all 5

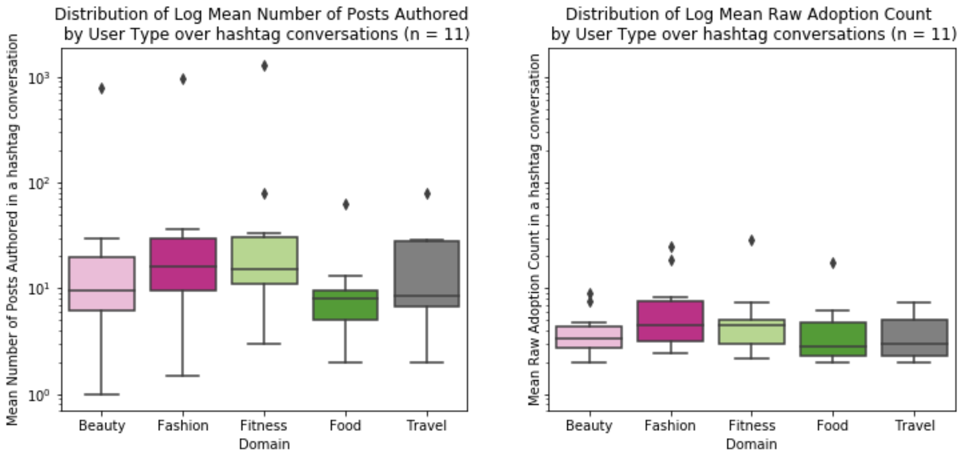

Figure 5.10: (Left) The distribution of mean number of posts authored by user type over hashtag conversations (Right) Distribution of mean raw adoption counts by user type over the subset of hashtag conversations

domains to minimise domain bias ($n = 10$). Here, we can see, for example, #nye has a significantly higher log adoption rate in comparison with other hashtag topics. #fitness, on the other hand, observes on average a lower log adoption rate, in comparison with the other hashtags, but not significantly lower than the second lowest: #ad ($t = 2.065, df = 7.0886, p = 0.0773$).

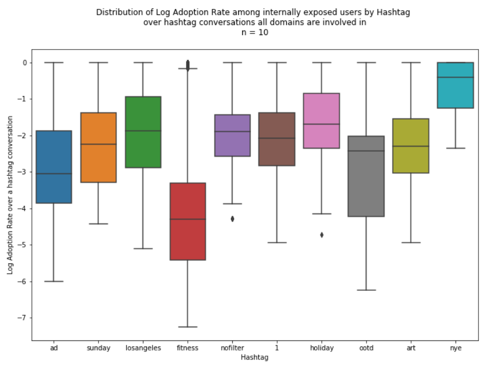

Figure 5.11: Distribution of log adoption rate over hashtag topic, in which influencers across all 5 domains have participated

Figure 5.12 shows the distribution of log adoption rate over domain and hashtag topic. This confirms that the variation in log adoption rate is observed by variations over both domain and hashtag topic ($F[4, 46443] = 7842.3, p < 2.2e^{-16}$ and $F[10, 46443] = 707.2, p < 2.2e^{-16}$ respectively).

Surprisingly, for hashtag topic #fitness, we see that fitness domain influencers have,

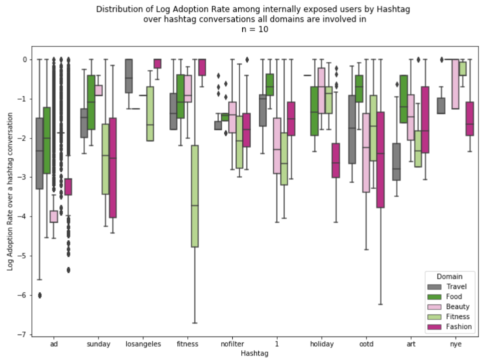

Figure 5.12: Distribution of log adoption rate over Domain and Hashtag Topic, in which influencers across all 5 domains have participated

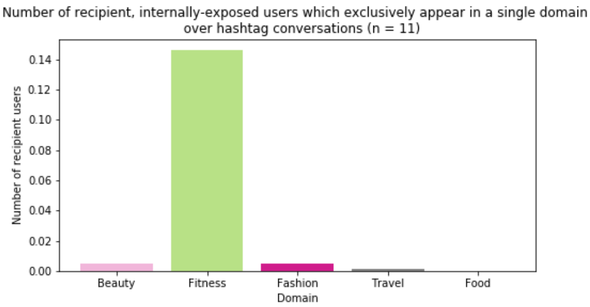

Figure 5.13: Number of internally-exposed readers, who are involved explicitly within group networks of influencers of a given domain

on average, a lower log adoption rate in comparison with other domains.

Exploring this further, Figure 5.13 shows the proportion of total readers involved over the 10 hashtags (n = 13977), which respond exclusively to posts from a single influencer domain (we will call these domain-specific readers). Among the domains, it appears that there is a distinctly large proportion of readers who exclusively respond to posts from fitness influencers ($\approx 14\%$) for the 10 common hashtags conversations.

Figure 5.14 shows that the majority of readers respond to posts from influencers across multiple domains ($> 75\%$) for the 10 common hashtags, with readers most commonly active in three domains.

Figure 5.15 shows the variation in log adoption rate for internally-exposed readers by domain for the common hashtags. We particularly note that the log adoption rate is highest among readers who either respond to posts from a single (1: ie. domain-specific) or from all (5: ie. domain-generic) influencer domains.

In summary, there is a significant variance in log adoption rates for domains and hashtag topic. We also witness a potential negative correlation between the log adop-

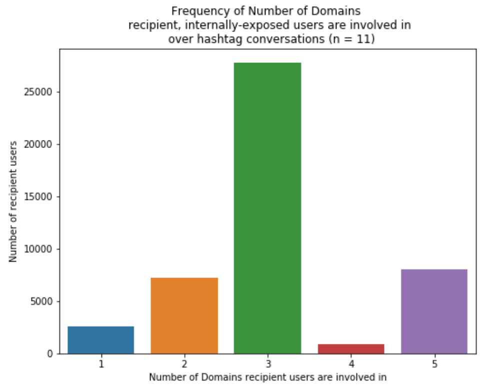

Figure 5.14: Frequency of internally-exposed readers by number of domains, in which the users are involved

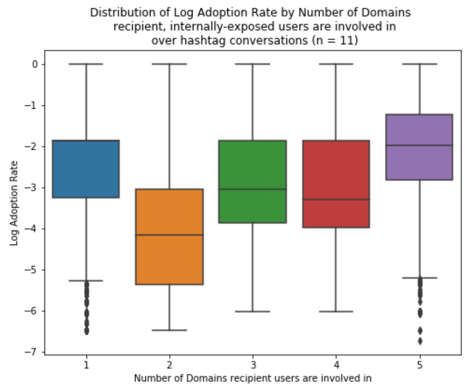

Figure 5.15: Distribution of log adoption rate by number of domains, in which internally-exposed readers are involved

tion rate and its parameters: the higher the log adoption rate, the less the influencer posts in a hashtag conversation and the lower the raw adoption count on average.

Furthermore, after exploring the effect of domain on the log adoption rate, we saw that, over the subset of hashtag conversations, there are very few domain-specific readers ($> 75\%$). The majority of readers particularly respond to influencers from 3 domains. Despite the low proportion of domain-specific readers, they still show a similar log adoption rate to those domain-generic readers who respond to influencers across all 5 domains.

Hypothesis 1.5 - influencers have greater authority over readers who exclusively respond to posts from a single domain

From Figure 5.16, when looking over all internally-exposed readers across all hashtag conversations, we see that the majority of readers are domain-specific ($> 90\%$).

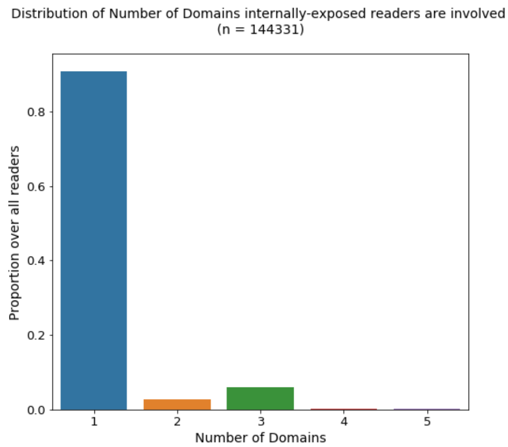

Figure 5.16: Distribution of internally-exposed readers by the number of domains they respond to over all hashtag conversations

Of those readers exclusive to a specific single domain, Figure 5.17 shows that the beauty domain has the most internally-exposed readers, with fitness and travel domains having the least.

From Figure 5.18, we observe significant variability in log adoption rate over the number of domains readers responds to ($F[4, 144326] = 238.7, p < 2e^{-16}$). It is notable that domain-generic readers show significantly higher log adoption rate than domain-specific (the next highest adoption rate) ($t = 14.491, df = 226.33, p < 2.2e^{-16}$).

Figure 5.19 shows, for domain-generic readers, every domain has relatively equal log adoption rate on average. However, for domain-specific show greater variation across the domains. The variation in log adoption can be explained by whether a reader responds to one domain or all five domains ($F[1, 131293] = 30.01, p = 4.31e^{-8}$). In particular, we see that readers exclusive to fashion influencers have a significantly greater average log adoption rate comparison to other domain influencers ($t = -39.386, df = 6432, p < 2.2e^{-16}$).

In summary, across all hashtag conversations, the vast majority of readers are domain-specific (i.e. exclusively respond to influencers from a single domain). We observe that the variation in log adoption rate can be explained by the number of domains

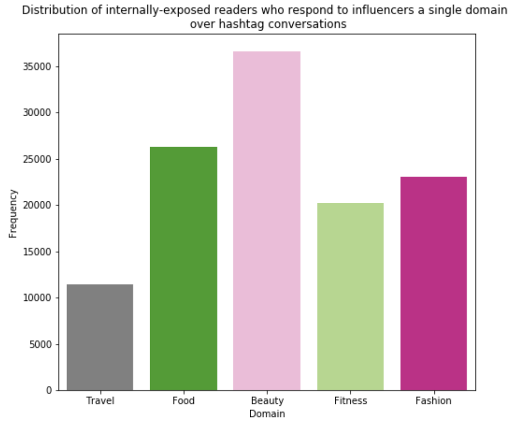

Figure 5.17: Distribution of internally-exposed readers, who respond to posts from influencers from a single domain

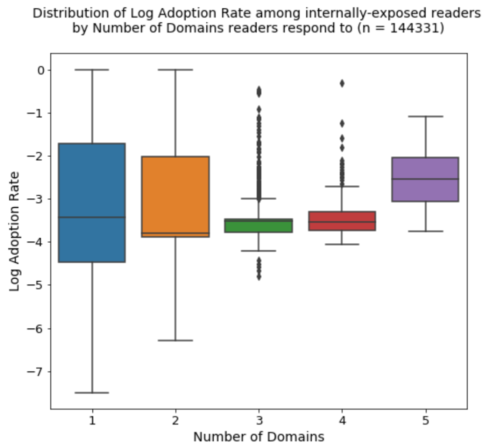

Figure 5.18: Distribution of log adoption rate among internally-exposed readers by the number of domains they respond to over all hashtag conversations

readers respond to. In common with Hypothesis 4.4, We observed that domain-specific readers present higher log adoption rate, on average, than readers who respond to between 2 and 4 domains. It is still the case that log adoption rate is the highest among domain-generic readers (i.e. those who respond to all 5 domains).

When comparing log adoption rates over domains by domain-specific and domain-generic readers, the log adoption rate is noticeably higher among domain-generic than among domain-specific readers. Also, in the majority of domains, the log adoption rate is consistent across all domains among domain-generic readers. with the exception of Fashion. The variation in log adoption can be explained by whether a reader responds to one domain or all five domains.

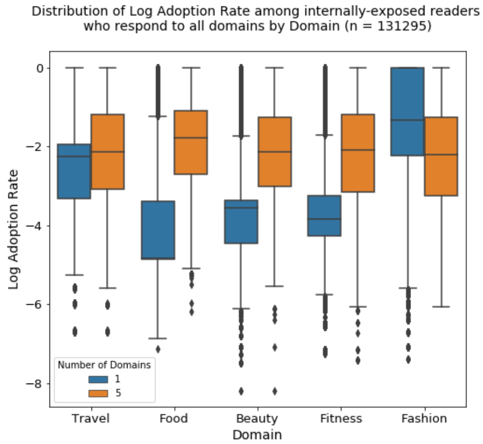

Figure 5.19: Distribution of log adoption rate by readers who respond to either a single, specific domain and those who respond to all five domains.

5.2 Persuasion

Hypothesis 2.1 - Influencers are more persuasive in hashtag conversations than regular users

Figure 5.20 shows that influencers in general have a noticeably greater variation in average user stylistic accommodation than regular users over hashtag conversations. We see that, for influencers, the distribution of user stylistic accommodation is significantly skewed towards zero (homophily) than among regular users ($t = -61.744, df = 91147, p < 2.2e^{-16}$). The regular users skew towards positive user stylistic accommodation.

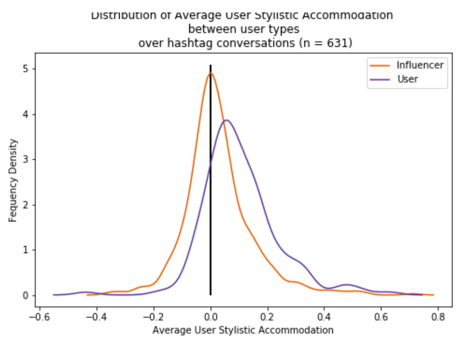

Figure 5.20: Distribution of average user stylistic accommodation over hashtag conversations across all hashtag conversations

From Figure 5.21, (left) we can see that for influencers the extent of the reader's accommodation is higher, on average, than for regular users ($t = 21.407, df = 83815, p < 2.2e^{-16}$). The author's extent of accommodation (middle), for the influencer, shows significantly greater change in linguistic style than for regular users ($t = 5.1705, df = 58014, p = 1.171e^{-7}$).

In regards to symmetry of accommodation (right), influencers are closer to zero in comparison with regular users.

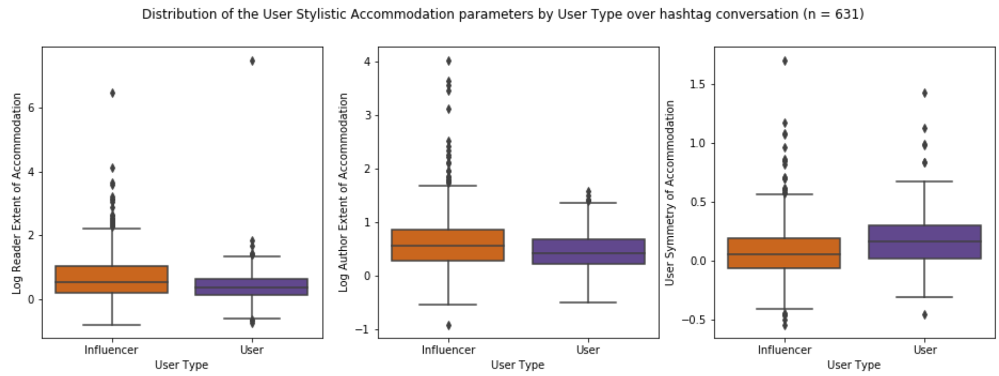

Figure 5.21: Distribution of reader extent (left), author (middle), user symmetry of accommodation (right), which are the parameters of the user stylistic accommodation function by user type

In summary, we see overall the influencer's distribution of User Stylistic Accommodation is skewed towards 0 in comparison with regular users. We also observe that influencers have significantly higher author and reader Extents of Accommodation than regular users, on average. Furthermore, we see that influencer's Symmetry of Accommodation is significantly lower and closer to 0 than regular users.

Hypothesis 2.2 - Domain and topic have significant effect over influencer's user stylistic accommodation with users

As we did in Hypothesis 1.2, we will segment situational influence to the domain- and topic-level. Here, we will assess whether the variation in user stylistic accommodation can be explained by the variation in influencer domain and hashtag conversation across hashtag conversation, in which influencers from all five domains participated (n = 11).

Figure 5.22 shows that all influencers have a user stylistic accommodation above 0 (ie. users are accommodating to the influencers style). The food and travel domains show the greatest user stylistic accommodation on average (≈ 0.65). Beauty and fitness influencer domains show similar distribution of average stylistic accommodation.

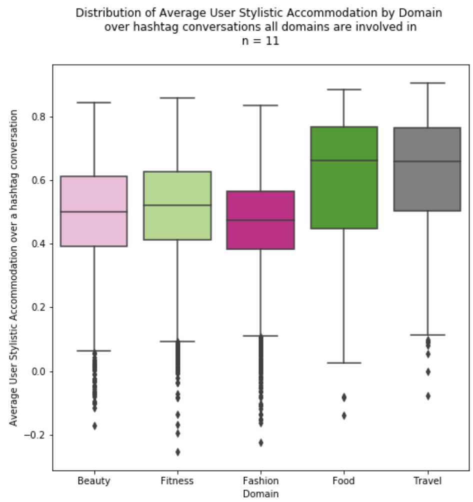

Figure 5.22: Distribution of user stylistic accommodation by domain over hashtag conversations, in which influencers from all 5 domains have participated

Fashion influencers, on average, show a lower user stylistic accommodation in comparison with other domains.

From Figure 5.23 shows user stylistic accommodation at the hashtag topic level (for the hashtags common to all 5 domains). Again all user accommodation is greater than zero. At the two extremes, we see that, influencers involved in the #ad topic, on average, experience higher user stylistic accommodation (≈ 0.7), while for the #momlife topic, there is little user stylistic accommodation (≈ 0.15).

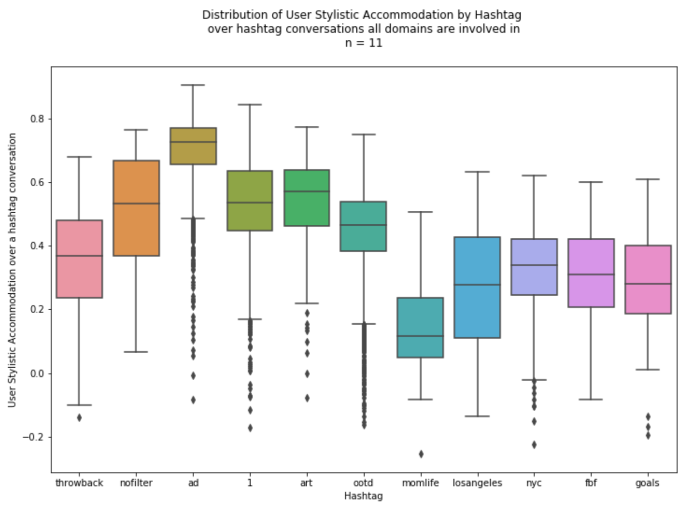

Figure 5.23: Distribution of user stylistic accommodation by hashtag over hashtag conversations, in which influencers from all 5 domains have participated

When we view the variation of user stylistic accommodation over domain and hashtag topic (Figure 5.24), we notice that some hashtag conversations show significant in mean user stylistic accommodation over domains (e.g. #throwback, #art) $F[4,10013] = 228.1, p < 2e^{-16}$. Variation in average User Stylistic Accommodation is also explained across hashtags $F[10,10013] = 624.4, p < 2e^{-16}$. There are also hashtag topics where average user stylistic accommodation varies over the domain (e.g. #losangeles, #momlife). This cross-domain variation tends to appear from hashtag topics which have lower mean user stylistic accommodation over hashtags.

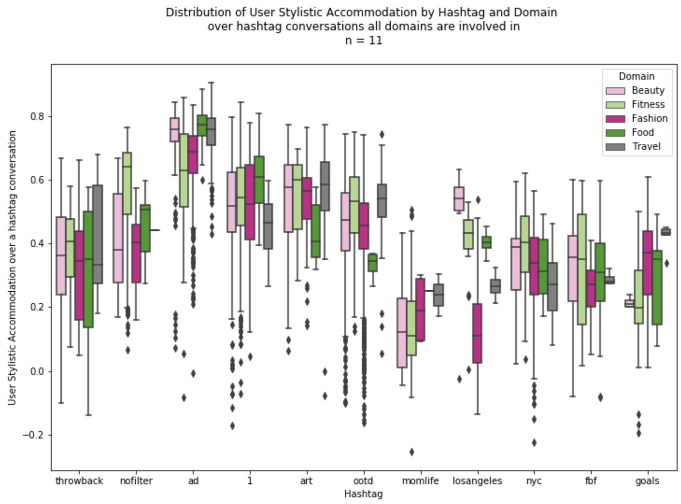

Figure 5.24: Distribution of user stylistic accommodation by hashtag and domain over hashtag conversations, in which influencers from all 5 domains have participated

When we review the distributions of the parameters of the user stylistic accommodation function (reader's extent [Figure 5.25], author's extent [Figure 5.26] and symmetry of accommodation [Figure 5.27]), we see that overall there is very little variation in these parameters across the domains.

In regards to the distribution of reader's extents of accommodation, we see marginal variation over topics and minimal variation over domains. It is worth noting the variation over reader's extent of accommodation over #losangeles. Fashion influencers witness significantly greater reader's extent in comparison with other domains ($t = 5.5083, df = 50.771, p = 6.045e^{-7}$, for example travel).

We also notice overall that the author's extent of accommodation hardly varies across domains and hashtag topics (See Figure 5.26).

From Figure 5.27), we observe overall variation in symmetry of change over topics, but not particularly across domains. #losangeles and #momslife topics show the least

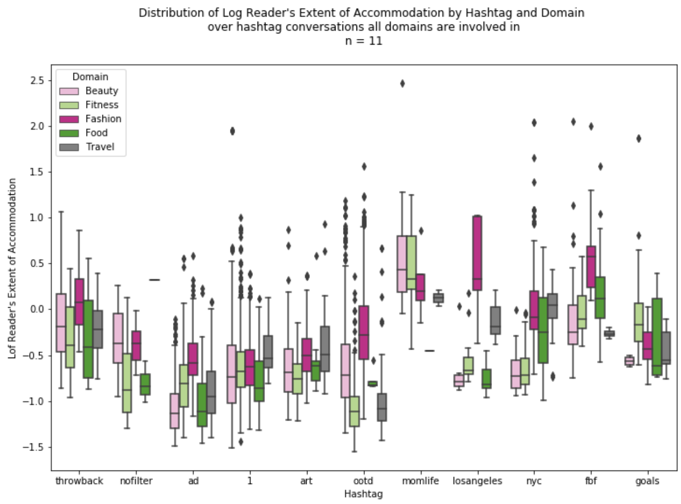

Figure 5.25: Distribution of log reader's extent of accommodation by domain and hashtag topic

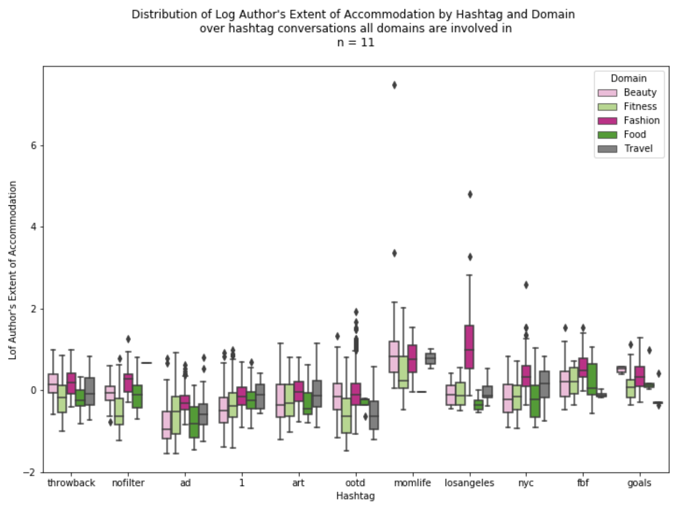

Figure 5.26: Distribution of log author's extent of accommodation by domain and hashtag topic

symmetry. It is also worth noting the high asymmetry of accommodation for the #ad topic.

In summary, over a subset of hashtag conversations, we saw that both domain and hashtags have a significant main effect, representing a variation in User Stylistic Accommodation. After reviewing the distribution of the parameters involved in the User Stylistic Accommodation function, overall, there is very little variation in author's Extent of Accommodation across hashtags and domains (except for notable examples

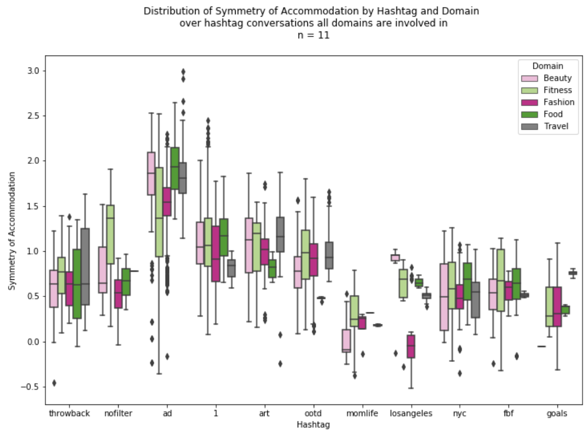

Figure 5.27: Distribution of symmetry of accommodation by domain and hashtag topic

such as #losangeles). We observe the majority of the variation occurring over the reader's Extent of Accommodation and the overall Symmetry of Accommodation both over domains and hashtags.

Hypothesis 3.1 - influencers have closer proximity to the hashtag norm than regular users

To observe whether influencers are closer to the hashtag norm than regular users, we select hashtag conversations in which both the regular users and influencers participate (n = 136).

To measure proximity between the regular user's and influencer's baseline styles and the hashtag norm, we normalised the baseline styles and hashtag norm, and applied cosine similarity between the normalised baseline styles and hashtag norms. We used normalised linguistic styles because we are concerned with the directionality than the magnitude.

Figure 5.28 shows that regular users are on average closer to hashtag topic norm than influencers.

Figure 5.29 shows the computed ratio between average cosine similarity over influencers and users by hashtag topic. A value of 1 suggests that influencers are as close to the hashtag norm as the user. A value greater than 1 suggests that influencers are closer to the hashtag norm than regular user and vice versa. The distribution is centred around 1 with a very low standard deviation (0.18).

Figure 5.30 shows the overall average cosine similarity between normalised influencer's and regular user's baseline linguistic style over all hashtag conversation. Typi-

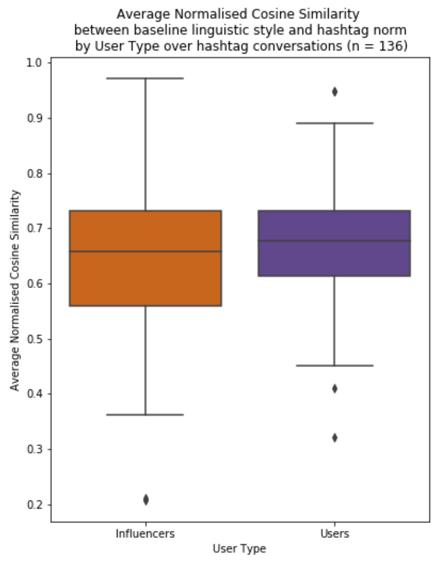

Figure 5.28: Distribution of cosine similarity between normalised author's baseline style and the hashtag norms over user type

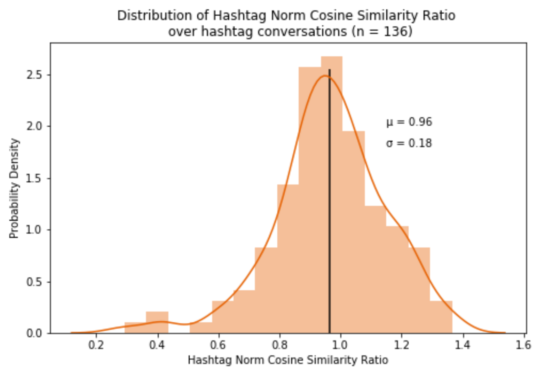

Figure 5.29: Distribution of cosine similarity ratio between influencers and users and the hashtag norm over hashtag conversations

cally, cosine similarity between influencers is between 0.4 and 0.8, on average over all hashtag conversations.

In summary, there is no significant difference in proximity to the hashtag norm for influencers and regular users. Further analysis shows that they are both equally close to the hashtag norm. We also saw that influencer's and regular user's baseline styles are relatively similar to each other (0.6 ± 0.2).

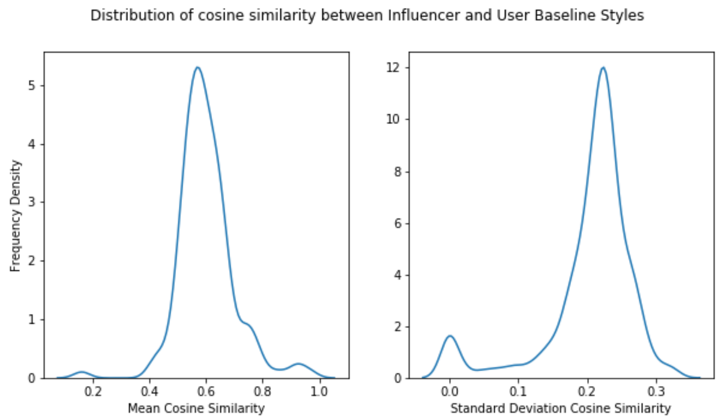

Figure 5.30: Distribution of mean (left) and standard deviation (right) over cosine similarity between normalised influencer and reader baseline styles

Hypothesis 3.2 - Association between Hashtag Norm proximity and Hashtag Stylistic Accommodation

Overall, we do see a positive relationship between hashtag norm proximity and hashtag stylistic accommodation over influencers. However, Figure 5.31 shows this relationship to be extremely weak ($R^2 = 0.103, \beta = 0.489$).

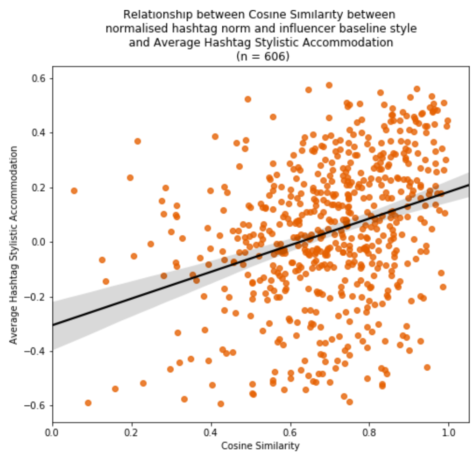

Figure 5.31: Weak positive relationship between hashtag norm proximity and hashtag stylistic accommodation by user type.

Hypothesis 4 - Influencers predominately use the peripheral (superficial) persuasion route over central (logic and reasoning) route

From Figure 5.32 we see that the Mean Attitude Change ranges between 0 and 0.238. The distribution is heavily skewed towards 0.

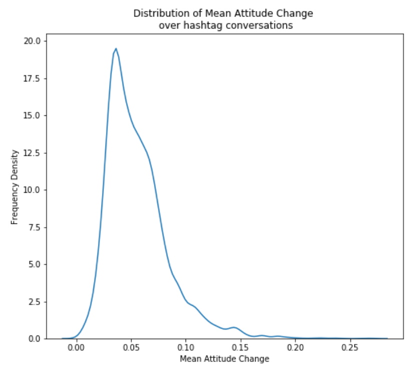

Figure 5.32: Distribution of Mean Attitude Change over hashtag conversations

Figure 5.33 demonstrates the lack of variation in mean attitude change by type and scale of user stylistic accommodation. However, we do see a pattern, whereby the average mean attitude change increases over the types of accommodation (i.e. Over homophily, mean attitude change is, on average, higher than $A \rightarrow R$ accommodation). $R \rightarrow A$ accommodation presents the highest the mean attitude change on average over the other types of accommodation.

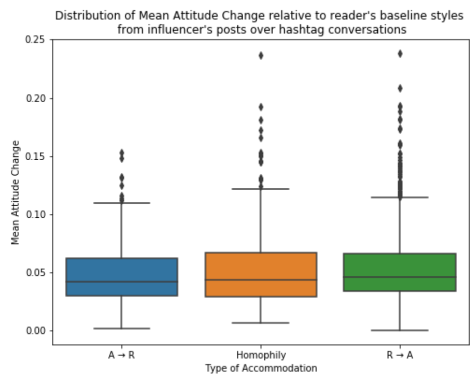

Figure 5.33: Distribution of Mean Attitude Change over hashtag conversations by type and scale of User Stylistic Accommodation. $A \rightarrow R$ accommodation : $(-0.9, -0.05]$, Homophily: $(-0.05, 0.05]$, $R \rightarrow A$ accommodation : $(0.05, 0.9]$

In summary, the distribution of Mean Attitude Change is skewed towards 0. When viewing the distribution of Mean Attitude Change over the types of accommodation, we see that as the reader moves more towards the author (from $A \rightarrow R$, through Homophily, to $R \rightarrow A$), the Mean Attitude Change increases, on average.

5.3 Topic Initiation and Uptake

Hypothesis 5.1: Influencers initiate more trends than regular users

We expect that influencers have a greater proportion of $TD > 0.6$ than users to signify trend setting.

Figure 5.34 shows that influencer and regular user topic dominance distributions are relatively similar. Influencers and regular users show peaks at $TD_t^A = 0$ and $TD_t^A = 1$. However, we do notice a noticeable bump for influencers at $TD_t^A = 0.5$.

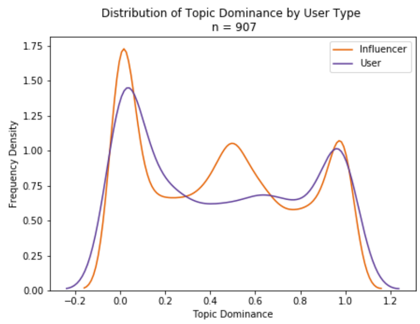

Figure 5.34: Distribution of Topic Dominance by user type over hashtag conversations

Figure 5.35 shows the frequency of Topic Dominance categories: $0 \geqslant TD_t^A > 0.4$, $0.4 \geqslant TD_t^A \geqslant 0.6$ and $TD_t^A > 0.6$ by user type. Here, we can see that influencers have almost double the proportion of $0.4 \geqslant TD_t^A \geqslant 0.6$ than regular users. We furthermore observe that regular users have a greater proportion of $0 \geqslant TD_t^A > 0.4$ and $TD_t^A > 0.6$ than influencers.

From Figure 5.36, we see that regular users have significantly greater log uptake gradients P_1' (left) and P_2' (right) than influencers on average ($t = -33.2, df = 3719.3, p < 2.2e^{-16}$ and $t = -33.255, df = 3713.2, p < 2.2e^{-16}$ for P_1' and P_2' respectively). This trend is also consistent over the Topic Dominance categories (See Figure 5.37).

We also notice a strong linear correlation between the log uptake gradients P'_1 and P'_2, which occurs over both user types ($R^2 = 0.935, \beta = 0.966$).

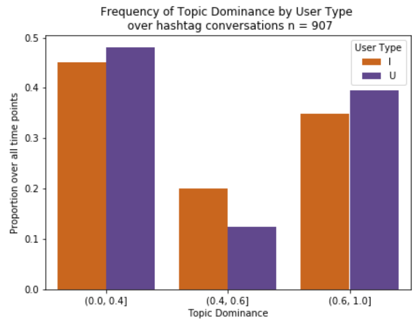

Figure 5.35: Frequency of Topic Dominance categories $0 \geqslant TD_t^A > 0.4$, $0.4 \geqslant TD_t^A \geqslant 0.6$ and $TD_t^A > 0.6$ by user type

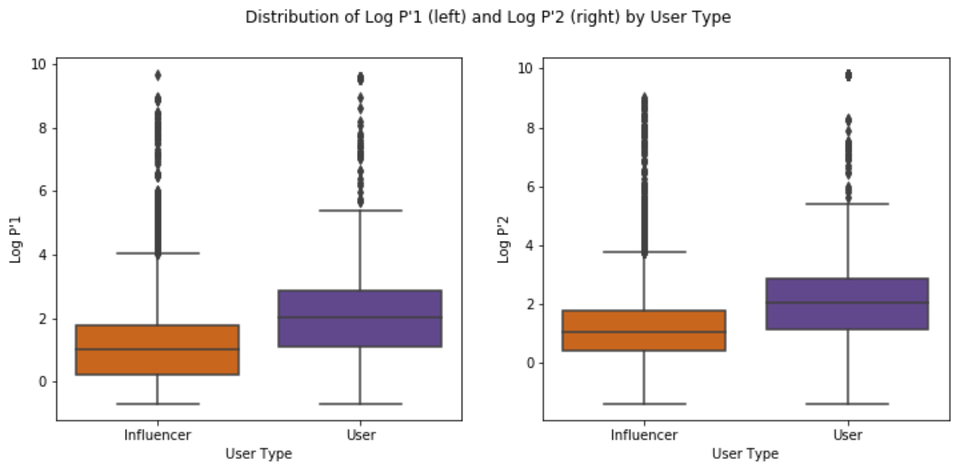

Figure 5.36: Distribution of log uptake gradients P_1' (left) and P_2' (right) by User Type

In summary, influencers have noticeably more time points, at which $0.4 \geqslant TD_t^A \geqslant 0.6$ than regular users. Furthermore, regular users do show a higher proportion of time points where $TD_t^A > 0.6$ than influencers.

Regular users have significantly higher uptake gradients (P_1' and P_2'), on average, than influencers. This trend was also consistent across Topic Dominance categories.

Hypothesis 5.2: Influencer trend initiation is domain- and topic-dependent

From Figure 5.39, we see that there are differing distributions of topic dominance by domain. For example, among fashion influencers we see that the majority of times when they post, topic dominance is zero ($TD_t^A = 0$). Among food influencers on the other hand, for the vast majority of times when they post, topic dominance is 0.5

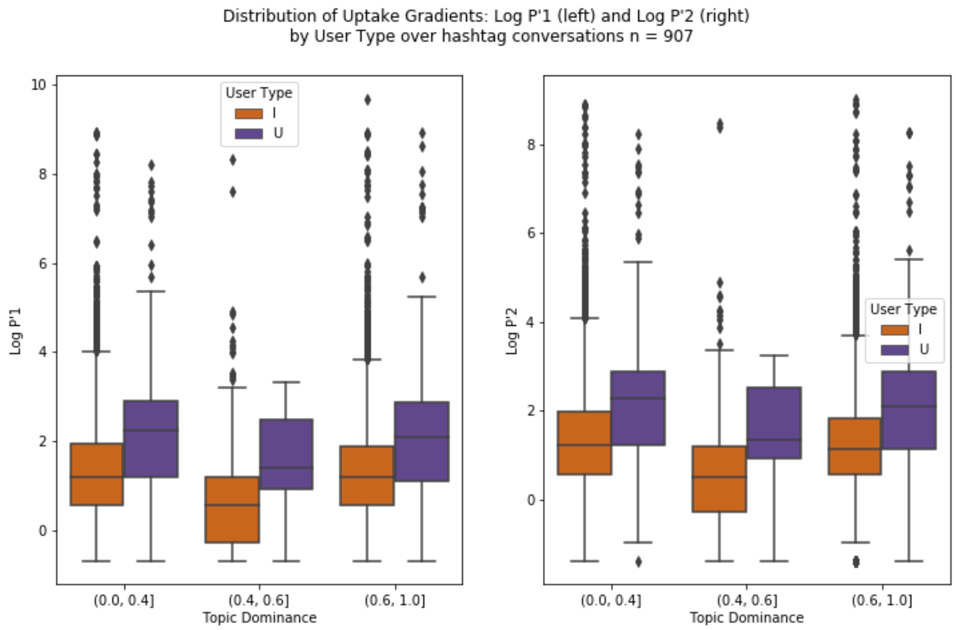

Figure 5.37: Distribution of Uptake Gradients: Log P'1 (left) and Log P'2 (right) by User Type over Topic Dominance categories over hashtag conversations

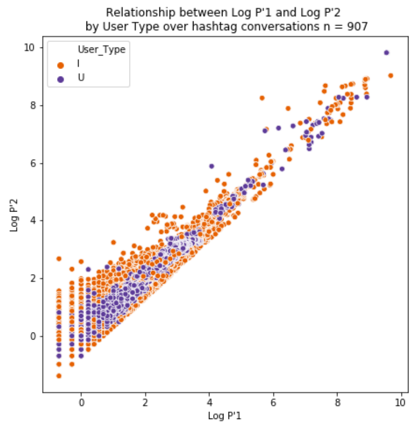

Figure 5.38: Positive relationship between log uptake gradients P_1' and P_2' by User Type

$(TD_t^A = 0.5)$.

Figure 5.40 further demonstrates the variation in topic dominance over domains, which is relatively small for $0 \geqslant TD_t^A > 0.4$ and $TD_t^A > 0.6$. the majority of variation can be observed in $0.4 \geqslant TD_t^A \geqslant 0.6$.

From Figure 5.41, we see an overall significant variation in log uptake gradients P_1' (left) and P_2' (right) over domains $(F[4, 15752] = 38.95, p < 2e^{-16})$.

It seems like there could be a relationship between the distribution of uptake gradi-

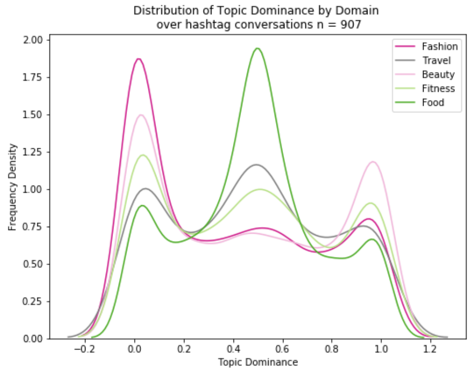

Figure 5.39: Distribution of Topic Dominance over Domain over all hashtag conversations

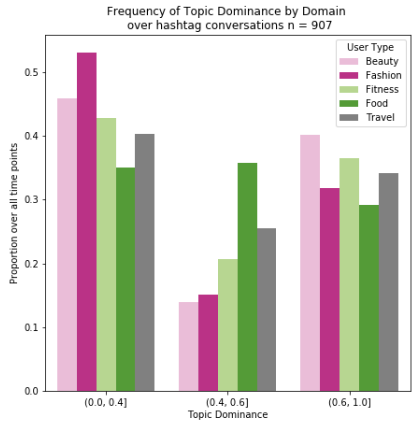

Figure 5.40: Frequency of Topic Dominance categories $0 \geqslant TD_t^A > 0.4$, $0.4 \geqslant TD_t^A \geqslant 0.6$ and $TD_t^A > 0.6$ by Domain

ents over domains and the distribution of Topic Dominance categories over domains. For example, for domains with lower proportions of $0.4 \geqslant TD_t^A \geqslant 0.6$ (e.g. Fashion and Beauty tend to have higher log uptake gradients on average than those which have more uniform distribution of Topic Dominance categories (e.g. Food) (See Figure 5.42).

In summary, there is notable variation in distribution of Topic Dominance over domains. We see that the majority of this variation particularly occurs at $0.4 \geqslant TD_t^A \geqslant 0.6$. We also noted a negative association between the proportion of $0.4 \geqslant TD_t^A \geqslant 0.6$ in a domain and log uptake gradients (P_1' and P_2').

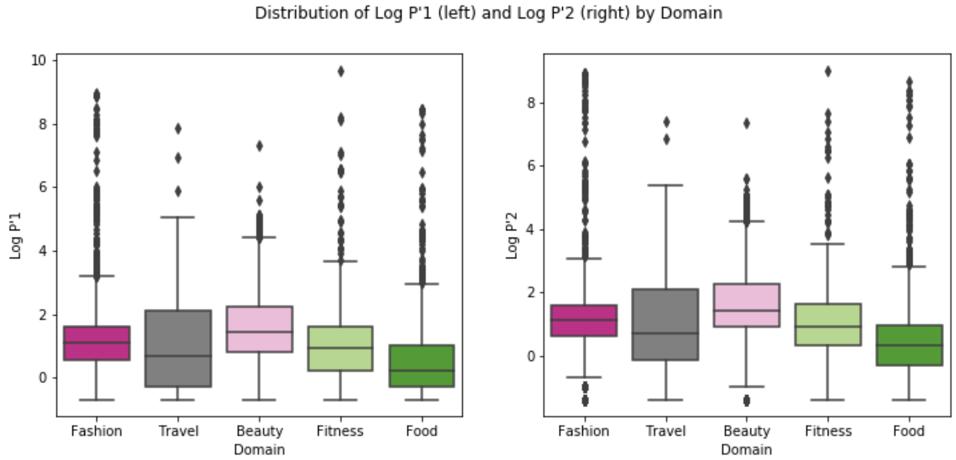

Figure 5.41: Distribution of log uptake gradients P_1' (left) and P_2' (right) by Domain

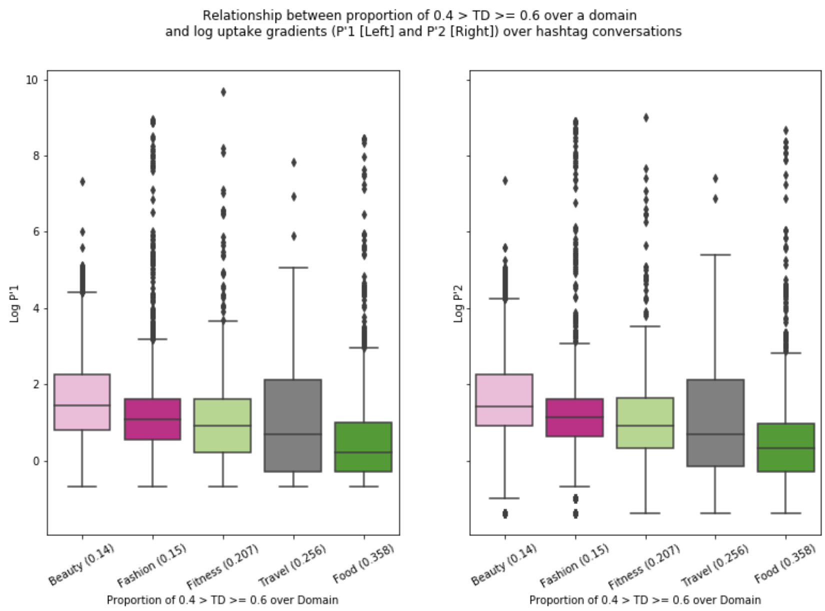

Figure 5.42: Distribution of log uptake gradients (P_1' [left] and P_2' [right]) by proportion of $0.4 \geqslant TD_t^A \geqslant 0.6$ over domains.

Chapter 6

Discussion

6.1 Authority

Overall, the findings from Hypotheses 1.1 - 1.5, show an overall paradox in the authority dimension.

The importance of the reader in idea adoption is intrinsic in the definition of authority: The rate of adoption over a hashtag conversation for a given author is dependent on how many readers respond to their post within the timeframe. The results revealed that the influencer's authority over a domain is dependent on the domain-specificity of their reader (number of domains a reader responds to). This aspect fall in line with Watts and Dodds, 2007's notion of "accidental influentials".

However, the noticeable higher adoption rates of influencers compared with regular users shows that influencers do have higher authority.

We also noticed that influencers had a higher proportion of internally-exposed users (i.e. readers who frequently interact with the author) and therefore reinforces the idea that influencers had a greater level of implicit connectivity between them and their readers compared to regular users, and thus provides more opportunities for influencers to diffuse ideas over more readers. Similar proportions of externally- and internally-exposed readers were found in Myers et al., 2012's formulation measuring diffusion of links on Twitter.

On the global level, we found that an influencer's authority can be associated with two main factors:

1. High scale of implicit connectivity: This means that the post reaches a larger group of readers, as measured by the author's group network size. In the case of influencers, they show significantly greater scale of implicit connectivity than regular users (Hypothesis 1.2).

2. High strength of implicit connectivity: This means that the author is retaining a larger proportion of their audience, as measured by the group network size / audience proportion. The influencers show higher strength of implicit connectivity than regular users, irrespective of the audience size (Hypothesis 1.2).

It is surprising to note that this difference in authority over user types is not explained by the low raw adoption counts for influencers. Also, despite the differences not being significant, we do observe influencers post less than regular users on average. This suggests that authority is not brought about by posting more in a hashtag conversation. It also suggests that authority favours more the scale of implicit connectivity with readers, than the strength.

This contradicts findings from Moldovan et al., 2017, which suggests that overall greater authority could be a result of influencers leveraging smaller scale, strongly-connected groups.

Overall, authority in the context of social media, appears to result from reaching a wide number of unique readers (high scale) with the fewest number of posts. The number of times that the reader adopts the hashtag is of little consequence. Hence why we observe influencers having significantly greater group network sizes over regular users.

Unlike at the global level, both higher scale and strength of implicit connectivity does not translate into higher authority at the domain level.

For example, in Hypothesis 1.3, we saw that Fashion influencers had the highest proportion of internally-exposed users, which can be explained by the almost consistently high scale and strength of implicit connectivity, regardless of the audience size. However, we observed that fashion influencers had the second lowest log adoption rate, on average over all domains.

However, we did notice higher log adoption rates are associated with lower number of posts and, surprisingly, lower raw adoption count. Food influencers, who experience a higher log adoption on average, tend to post less in a hashtag conversation on average. They also have readers who less frequently adopt the hashtag after them on average.

This reinforces the idea that greater authority is based reaching a larger group of readers with fewer posts.

Sitting under the significant effect of domain on the log adoption rate, we also saw that the domain-specificity of the reader is a key factor.

As we saw in Hypothesis 1.4 and 1.5, authority was highest over domain-generic readers (i.e. those who respond to influencers from all 5 domains), followed closely by domain-specific readers (i.e. those who exclusively respond to influencers from a single domain).

From Myers et al., 2012 model of authority, it is expected that the increase in exposure to the hashtag increases the probability that the hashtag will be adopted by the reader. This would explain why domain-generic readers have a higher log adoption rate. However, it does not explain why log adoption rate over domain-specific readers is greater than those readers who respond to influencers over 2, 3 & 4 domains.

A possible explanation for this can be shown in Figure 5.18: The log adoption rate for Fashion-specific readers is higher than other domains and thus increases the domain-specific log adoption rate, on average. For Food, Travel, Beauty and Fitness influencers have lower authority over domain-specific than domain-generic readers.

Hence, it is only Fashion influencers which notice a benefit to authority over a more domain-specific audience. This is in line with intuition, in that specific fashion readers will be looking exclusively to fashion influencers for trends.

With the exception of Fashion, influencers have noticeably greater authority over domain-generic readers than domain-specific readers.

In addition, Figure 5.19 revealed that all domains exert very consistent authority over domain-generic readers that respond to posts from all five domains. This suggests that all domains equally benefit over domain-generic readers.

The results have also shown support for Rosenthal and McKeown, 2016's notion that influencers typically are influential over particular situations (i.e. over a domain and or topic).

As well as seeing a main effect in domain, we also saw a main effect in topic. This suggests therefore that the degree of authority an author has over an audience depends on the domain (including the domain-specificity of the reader), as well as the topic.

6.2 Persuasion

Across all levels of persuasion, we see that the degree of persuasion involves more of a collaboration between influencer and reader.

From Hypothesis 2.1, we saw that, on the global level, influencers significantly skew more towards homophily than regular users (this skew we will denote as a ”community effect”).

Therefore, the question we should ask in this dimension of Influence Space is not one of how persuasive an influencer is, but to what extent do influencers align with their community of readers. Rosenthal and McKeown, 2016 even goes so far to say that by achieving homophily with the audience, an influencer can be more influential.

From the reader's perspective, we saw that influencer's readers change their attitude (higher reader extent of accommodation) more than readers of regular users (Figure 5.21 left)).

In addition, from the author's perspective, we saw that influencers change their attitude (higher author extent of accommodation) more than regular users (Figure 5.21 middle)). These extents of accommodation from both author and reader, in combination with the symmetry of accommodation indicate that influencers and their readers are symmetrically converging toward each other. This evidences persuasion (specifically homophily) as collaboration between influencer and reader. This collaboration follows the principles from Šćepanović et al., 2017 where both influencers and readers are changing their attitudes so that they are more similar to each other.

We must highlight that regular users were significantly more persuasive than influencers. By considering regular users as a proxy for an influencer's audience, we could infer that the influencer's immediate audience is the source of their persuasive power. The combination of influencers achieving similarity, to appeal to their audience, and regular users showing stronger persuasion over their readers provides a collaboration where second-degree persuasion can occur.

In Hypothesis 4, we evaluated the persuasion route which influencers implement, relative to the Elaboration Likelihood Model Petty and Cacioppo, 2012. The analysis provides evidence to suggest that influencers implement a more central route to persuasion.

The distribution of Mean Attitude Change was heavily skewed towards zero, with the majority of readers showing minimal fluctuation in their linguistic style after the influencer's persuasion attempt (0.05). This appears to be in line with Wood, 2000's core assumption that an attitude towards something is relatively stable over time.

Since the range of Mean Attitude Change is between $0 - 1.49$, where 0 indicates a central route and 1.49 indicates a peripheral route, then we can say that influencers are more likely implement a more central route to persuasion.

To see if persuasion strategy changes to achieve a type of accommodation, we reviewed the distribution of the Mean Attitude Change over the three types of accommodation ($A \rightarrow R$, Homophily and $R \rightarrow A$). Across all types of accommodation, Mean Attitude Change remains consistently low, which evidences that influencers consistently use a central route to persuasion. From ELM, the central route indicates that readers seek more information in order to change their attitude. In this case, this suggests that influencers are used as a reliable source of information (explained by significantly higher authority) and therefore, their attitude remains stable after the influencer's persuasion attempt.

Furthermore, at the domain level, we saw that the reader makes the majority of the effort in persuasion.

Firstly, The analysis showed very little variation in user stylistic accommodation over the domains and thus suggests that no particular set of influencers are particularly more persuasive.

It is, however, worth highlighting the very high cross-domain variation in #losangeles, for example. The variation in user stylistic accommodation over the domains is reflected in the variation of the reader's extent of accommodation, which again places more evidence to suggest that persuasion is more determined by the reader's response, than the influencer themselves. Furthermore, in regards to author's (i.e. influencer's) extent of accommodation, we see very little variation overall, across hashtag topics as well as domain. This suggests that the influencer is not deviating from their usual baseline linguistic style and thus remaining authentic.

On the topic-level, we see that persuasion does significantly vary over topic and thus further evidences Rosenthal and McKeown, 2016's idea of influencers possessing situational "influence".

The "community effect" is also observed at the hashtag topic level. Although influencer's baseline styles are not closer to hashtag norms than regular users, influencers and regular users are actually similar in several ways. Firstly, influencers and regular user are as close to hashtag norms as each other. Secondly, this can be partially explained by the fact that influencer and regular user's baseline styles are already similar to each other in the first instance, which confers the simplified view of homophily as similarity between users De Choudhury et al., 2010.

The topic-level community effect from influencers did not significantly translate into greater persuasion, which does not conform with findings within the literature (Rosenthal and McKeown, 2016, Cartwright, 1951, Platow et al., 2015, Turner, 1991). We did see a very weak positive relationship between hashtag norm proximity and hashtag stylistic accommodation where the gradient was minimal. This suggests that being closer to the group norm does not yield a noticeably large difference in persuasiveness over readers. The deviation from the findings could be brought about by the nature of the studies themselves, which primarily focus on user homophily as similarity in user traits, rather than the communication between users (as proposed by Shepard, 2001, Šćepanović et al., 2017). For example, Rosenthal and McKeown, 2016 measures user similarity by author traits (e.g. age, religion, etc.).

Overall, we observed a community effect at both the global and topic levels. However, at the domain level, influencers tended more to $R \leftarrow A$, instead of homophily. This could be explained by the properties of the hashtag conversations involved in the analysis, which can be explored by further analysis.

6.3 Topic Initiation and Uptake

On the global level, we noticed a key difference in the distributions of topic dominance between influencers and regular users: Influencers surprisingly have a greater tendency to follow trends then set trends in comparison with regular users.

Across both influencers and regular users, when setting new trends, they saw greater uptake after their posts ($P_2^{'}$) and over the time window ($P_1^{'}$) than when they follow trends. However, influencers on average see smaller uptakes both before and after their posts in comparison with than regular users.

Although, these findings contradict a traditional view of an influencer as one who consistently setting trends, these findings are in line with Biran et al., 2012's original definition of the topic initiation and uptake dimension. Part of Biran et al., 2012's definition mentions the influencer's ability to respond to and follow user's interests and trends. This, of course, requires influencers to respond to trends when they appear, and therefore 'ride a wave' and follow a trend set by users.

At the domain level, we saw that the proportion of times influencers follow trends ($0.4 \geqslant TD_t^A \geqslant 0.6$) noticeably varies over domains. For example, intuitively, Fashion and Beauty influencers are less likely to follow trends than influencers in other domains. This is because in the Fashion and Beauty domains, influencers are expected to set new trends. This shows that these domains are noticeably volatile, i.e. the majority of posts either do not set a trend ($0 \geqslant TD_t^A \geqslant 0.4$) or set a trend ($TD_t^A > 0.6$).

The Food domain is shown to be less volatile, since a Food influencer appears to be equally likely to either set a trend, follow one or not set one at all.

We have shown that more volatile domains (e.g. Fashion and Beauty) tend to see greater uptake after their posts ($P_2^{'}$) and within the overall time windows ($P_1^{'}$) than less volatile domains (e.g. Food). For example, Fashion and Beauty domains witness noticeably larger uptake gradients, in comparison with the Food domain. This suggests that for influencers in less volatile domains play it safe (equally likely to either follow or set trends) and, as a result, receive a smaller uptake and than those who are in more volatile domains. Taking a risk in volatile domains is more worthwhile since it results in greater uptakes.

Chapter 7

Conclusion

The goal of this report was to get closer to answering how Instagram influencers are influential. The analysis has provided a vital and necessary framework in which to evaluate each dimension of Influence Space, which has overall shown that particularly authority and persuasion involves reader-author collaboration: neither just the reader nor influencer is responsible for the influential behaviours, which do significantly vary in contrast with regular users.

Firstly, we saw reader-author collaboration does help determine the influencer's authority, which aligns with the existing literature on authority (particularly, Platow et al., 2015). The results do show that domain-specificity of the reader does contribute to an influencer's amount of authority. However, globally, the influencer is shown to retain significantly more of its audience, regardless of the audience size. However, over all of this, influencer's authority strategy appears to be focused to get as many users adopting a hashtag at least twice with a minimal number of posts as possible.

We also saw a similar idea of reader-author collaboration in determining the amount of persuasion an influencer has on its audience. For example, globally, influencers significantly tend to align with their community and display homophily with their readers. On the topic level, We also saw that regular user's and influencer's baseline styles are similar to each other, which abides by the more traditional definition of homophily. However, when it comes to persuasion itself, the influencers' readers tend to make most of the effort in changing their attitudes towards the influencer.

In regards to topic initiation and uptake, we observed the importance of the domain itself in determining how much uptake follows an influencer's contribution into a hashtag conversation. Essentially, more volatile domains (such as Fashion) tend to give a greater payback for trend-setting, in the form of higher uptakes. This payback is in exchange for the risk that influencers might also receive minimal/no uptake after their contribution. Less volatile domains (e.g. Food) have an equal chance of setting a trend, following one, or not setting a new one and experience lower uptakes after their contributions as a compromise.

These insights show the general behaviour of influencers according to a robust framework, which assess influence across dimensions and levels. However, to obtain explicit patterns in influential behaviour, there needs to be a more granular analysis. Furthermore, a natural extension to the framework would be to include images processing to accommodate for the visual nature of the Instagram social media platform. This report, however provides a necessary framework to conduct such an analysis.

In understanding these behaviours in greater detail, we can apply this knowledge into improving influencer detection, so that we can more confidently discriminate regular users from influencers. More importantly, it is heartening that this analysis shows influential behaviours are not solely dependent on the influencer. Therefore, as the reader, you also determine the fate of influencers and opinion leaders.

References

25 top foodie influencers. (2017). Izea. https://izea.com/2017/07/10/25-top-foodie-influencers/

Agam, D. (2017). Followers ratio on instagram affects in product's brand awareness. *Australian Journal Accounting, Economics, and Finance (AJAEF)*, *3*(2), 85–89.

Agrawal, H. (2019). 11 top fashion bloggers on instagram of 2019 [best instagram outfits]. esocmedia. https://www.esocmedia.com/influencer-marketing/11-top-fashion-bloggers-on-instagram-of-2019-best-instagram-outfits/

Anger, I., & Kittl, C. (2011). Measuring influence on twitter, In *Proceedings of the 11th international conference on knowledge management and knowledge technologies*.

Azuma, S. (1997). Speech accommodation and japanese emperor hirohito. *Discourse & Society*, *8*(2), 189–202.

Bakshy, E., Hofman, J. M., Mason, W. A., & Watts, D. J. (2011). Everyone's an influencer: Quantifying influence on twitter, In *Proceedings of the fourth acm international conference on web search and data mining*.

Bartone, P. T., Eid, J., Johnsen, B. H., Laberg, J. C., & Snook, S. A. (2009). Big five personality factors, hardiness, and social judgment as predictors of leader performance. *Leadership & Organization Development Journal*.

Beal, G. M., & Bohlen, J. M. (1956). *The diffusion process* (tech. rep.).

Beauty influencers: The top 25 you need to follows. (2017). Izea. https://izea.com/2017/05/25/top-beauty-influencers/

Biran, O., Rosenthal, S., Andreas, J., McKeown, K., & Rambow, O. (2012). Detecting influencers in written online conversations, In *Proceedings of the second workshop on language in social media*. Association for Computational Linguistics.

Bohlen, J. M. (1967). Needed research on adoption models1. *Sociologia Ruralis*, *7*(2), 113–129.

Cartwright, D. (1951). Achieving change in people: Some applications of group dynamics theory. *Human Relations*, *4*(4), 381–392.

Cha, M., Haddadi, H., Benevenuto, F., & Gummadi, K. P. (2010). Measuring user influence in twitter: The million follower fallacy, In *Fourth international aaai conference on weblogs and social media*.

Cialdini, R. B., & Goldstein, N. J. (2004). Social influence: Compliance and conformity. *Annu. Rev. Psychol.*, *55*, 591–621.

Coffman, L., & Niehaus, P. (2014). *Pathways of persuasion* (tech. rep.). Working paper.

Danescu-Niculescu-Mizil, C., Gamon, M., & Dumais, S. (2011). Mark my words!: Linguistic style accommodation in social media, In *Proceedings of the 20th international conference on world wide web*. ACM.

De Choudhury, M., Sundaram, H., John, A., Seligmann, D. D., & Kelliher, A. (2010). " birds of a feather": Does user homophily impact information diffusion in social media? *arXiv preprint arXiv:1006.1702*.

Flaxman, S., & Kassam, K. (n.d.). On# agony and# ecstasy: Potential and pitfalls of linguistic sentiment analysis.

Forsey, C. (2019). The ultimate list of instagram influencers in every industry (135 and counting!) Statista. https://blog.hubspot.com/marketing/instagram-influencers

Global social media ranking 2019. (2020). Statista. https://www.statista.com/statistics/272014/global-social-networks-ranked-by-number-of-users/

Han, M., & Li, Y. (2018). Influence analysis: A survey of the state-of-the-art. *Mathematical Foundations of Computing*, *1*(3), 201–253.

Hogg, M. A., & Reid, S. A. (2006). Social identity, self-categorization, and the communication of group norms. *Communication theory*, *16*(1), 7–30.

Jones, S., Cotterill, R., Dewdney, N., Muir, K., & Joinson, A. (2014). Finding zelig in text: A measure for normalising linguistic accommodation, In *Proceedings of coling 2014, the 25th international conference on computational linguistics: Technical papers*.

Judge, T. A., Bono, J. E., Ilies, R., & Gerhardt, M. W. (2002). Personality and leadership: A qualitative and quantitative review. *Journal of applied psychology*, *87*(4), 765.

Katz, E., & Lazarsfeld, P. F. (1966). *Personal influence, the part played by people in the flow of mass communications*. Transaction publishers.

Kusner, M., Sun, Y., Kolkin, N., & Weinberger, K. (2015). From word embeddings to document distances, In *International conference on machine learning*.

Le, Q., & Mikolov, T. (2014). Distributed representations of sentences and documents, In *International conference on machine learning*.

Mikolov, T., Sutskever, I., Chen, K., Corrado, G. S., & Dean, J. (2013). Distributed representations of words and phrases and their compositionality, In *Advances in neural information processing systems*.

Miller, G. R. (2013). On being persuaded: Some basic distinctions.

Moldovan, S., Muller, E., Richter, Y., & Yom-Tov, E. (2017). Opinion leadership in small groups. *International Journal of Research in Marketing*, *34*(2), 536–552.

Myers, S. A., Zhu, C., & Leskovec, J. (2012). Information diffusion and external influence in networks, In *Proceedings of the 18th acm sigkdd international conference on knowledge discovery and data mining*.

Nasir, M., Chakravarthula, S. N., Baucom, B., Atkins, D. C., Georgiou, P., & Narayanan, S. (2019). Modeling interpersonal linguistic coordination in conversations using word mover's distance. *arXiv preprint arXiv:1904.06002*.

Nguyen, V.-A., Boyd-Graber, J., Resnik, P., Cai, D. A., Midberry, J. E., & Wang, Y. (2014). Modeling topic control to detect influence in conversations using non-parametric topic models. *Machine Learning*, *95*(3), 381–421.

Numerical (second) derivative of time series data. (2018). Math.StackOverflow. https://math.stackexchange.com/questions/2875173/numerical-second-derivative-of-time-series-data

Pennebaker, J. W., & King, L. A. (1999). Linguistic styles: Language use as an individual difference. *Journal of personality and social psychology*, *77*(6), 1296.

Petty, R. E., & Cacioppo, J. T. (2012). *Communication and persuasion: Central and peripheral routes to attitude change*. Springer Science & Business Media.

Platow, M. J., Haslam, S. A., Reicher, S. D., & Steffens, N. K. (2015). There is no leadership if no-one follows: Why leadership is necessarily a group process. *International Coaching Psychology Review*, *10*(1), 20–37.

Platow, M. J., Van Knippenberg, D., Haslam, S. A., Van Knippenberg, B., & Spears, R. (2006). A special gift we bestow on you for being representative of us: Considering leader charisma from a self-categorization perspective. *British Journal of Social Psychology*, *45*(2), 303–320.

Reitter, D. (2017). Alignment in Web-based Dialogue: Who Aligns, and how Automatic is it? Studies in Big-Data Computational Psycholinguistics (M. N. Jones, Ed.). In M. N. Jones (Ed.), *Big Data in Cognitive Science*. Routledge Psychology Press.

Rosenthal, S., & McKeown, K. (2016). Social proof: The impact of author traits on influence detection, In *Proceedings of the first workshop on nlp and computational social science*.

Šćepanović, S., Mishkovski, I., Gonçalves, B., Nguyen, T. H., & Hui, P. (2017). Semantic homophily in online communication: Evidence from twitter. *Online Social Networks and Media*, *2*, 1–18.

Shen, C.-w., Kuo, C.-J., Ly, M., & Thi, P. (2017). Analysis of social media influencers and trends on online and mobile learning. *International Review of Research in Open and Distributed Learning*, *18*(1), 208–224.

Shepard, C. A. (2001). Communication accommodation theory. *The New Hand-book of Language and Social Psychology*, 33–56.

Sun, B., & Ng, V. T. (2012). Identifying influential users by their postings in social networks, In *Ubiquitous social media analysis*. Springer.

The 10 top travel instagram influencers whose journeys you should follow in 2019. (2019). mediakix. https://mediakix.com/blog/top-travel-instagram-influencers-best-popular/

The top fitness influencers to guide your workout. (2019). Izea. https://izea.com/2019/02/20/top-fitness-influencers/

Turner, J. C. (1991). *Social influence.* Thomson Brooks/Cole Publishing Co.

Valente, T. W., & Davis, R. L. (1999). Accelerating the diffusion of innovations using opinion leaders. *The Annals of the American Academy of Political and Social Science*, *566*(1), 55–67.

Vasilynonok, D. (2018). Danielv775/natural-language-identification-graduate-project. https://github.com/danielv775/Natural-Language-Identification-Graduate-Project

Watts, D. J., & Dodds, P. (2007). The accidental influentials. *Harvard Business Review*, *85*(2), 22–23.

Wood, W. (2000). Attitude change: Persuasion and social influence. *Annual review of psychology*, *51*(1), 539–570.